arXiv:1204.3508v2 [math.AG] 15 Jan 2013

LINEAR SERIES ON METRIZED COMPLEXES OF ALGEBRAIC CURVES

OMID AMINI AND MATTHEW BAKER

Abstract. A *metrized complex of algebraic curves* over an algebraically closed field κ is, roughly speaking, a finite metric graph Γ together with a collection of marked complete nonsingular algebraic curves C_v over κ, one for each vertex v of Γ; the marked points on C_v are in bijection with the edges of Γ incident to v. We define linear equivalence of divisors and establish a Riemann-Roch theorem for metrized complexes of curves which combines the classical Riemann-Roch theorem over κ with its graph-theoretic and tropical analogues from [AC, BN, GK, MZ], providing a common generalization of all of these results. For a complete nonsingular curve X defined over a non-Archimedean field $\mathbb{K}$, together with a semistable model $\mathfrak{X}$ for X over the valuation ring R of $\mathbb{K}$, we define a corresponding metrized complex $\mathfrak{C}\mathfrak{X}$ of curves over the residue field κ of $\mathbb{K}$ and a canonical specialization map $\tau_*^{\mathfrak{C}\mathfrak{X}}$ from divisors on X to divisors on $\mathfrak{C}\mathfrak{X}$ which preserves degrees and linear equivalence. We then establish generalizations of the specialization lemma from [B] and its weighted graph analogue from [AC], showing that the rank of a divisor cannot go down under specialization from X to $\mathfrak{C}\mathfrak{X}$. As an application, we establish a concrete link between specialization of divisors from curves to metrized complexes and the Eisenbud-Harris theory [EH] of limit linear series. Using this link, we formulate a generalization of the notion of limit linear series to curves which are not necessarily of compact type and prove, among other things, that any degeneration of a $\mathfrak{g}^r_d$ in a regular family of semistable curves is a limit $\mathfrak{g}^r_d$ on the special fiber.

Contents

1. Introduction 2
2. Metrized complexes of algebraic curves 10
3. A Riemann-Roch theorem for metrized complexes of curves 13
4. Background on Berkovich curves 21
5. Specialization inequality 28
6. Limit linear series 33
7. Example: Generic chains of genus-one curves 45
Appendix A. Rank-determining sets for metrized complexes 50
References 54

Acknowledgments: The authors would like to thank Vladimir Berkovich, Erwan Brugallé, Lucia Caporaso, Ethan Cotterill, Eric Katz, Johannes Nicaise, Joe Rabinoff, Frank-Olaf Schreyer, and David Zureick-Brown for helpful discussions and remarks. The second author was supported in part by NSF grant DMS-0901487.

1. Introduction

We begin by quoting the introduction from the groundbreaking paper [EH] of Eisenbud and Harris:

> "One of the most potent methods in the theory of (complex projective algebraic) curves and their linear systems since the work of Castelnuovo has been degeneration to singular curves... Most problems of interest about curves are, or can be, formulated in terms of (families of) linear series. Thus, in order to use degenerations to reducible curves for studying smooth curves, it is necessary to understand what happens to linear series in the course of such a degeneration, and in particular, to understand what structure on a reducible curve plays the part of a linear series."

Eisenbud and Harris were able to settle a number of longstanding open problems with their theory of limit linear series, and the theory has undergone numerous developments and improvements in the last 25 years. However, the theory of limit linear series only applies, for the most part, to a rather restricted class of reducible curves, namely those of *compact type* (i.e., nodal curves whose dual graph is a tree).

In [B], the second author introduced a new framework for degenerating linear series on curves, degenerating linear series on a regular semistable family of curves to a linear series on the *dual graph* of the special fiber. This theory, which is closely related to tropical geometry and also to the theory of Berkovich analytic spaces, is more or less orthogonal to the Eisenbud-Harris theory, in that it works best for special fibers which are *maximally degenerate*, meaning that the dual graph has first Betti number equal to the genus of the generic fiber. Specializing linear series to the dual graph provides no information whatsoever when the special fiber is of compact type. Intriguingly, both the Eisenbud-Harris theory and the second author's theory from [B] lead to simple proofs of the celebrated Brill-Noether theorem of Griffiths-Harris (see [CDPR] and Section 7 below for a comparison of the methods involved).

The aim of the present paper is to introduce a theoretical framework suitable for generalizing both the Eisenbud-Harris theory of limit linear series and the second author's theory of specialization from curves to graphs. The main new object of study is what we call a *metrized complex of algebraic curves* over an algebraically closed field κ; this is, roughly speaking, a finite metric graph Γ together with a collection of marked complete nonsingular algebraic curves C_v over κ, one for each vertex v (with a marked point on C_v for each edge of Γ incident to v). We define divisors, linear equivalence of divisors, and rank of divisors on metrized complexes of curves and establish a Riemann-Roch theorem for metrized complexes of curves which combines the classical Riemann-Roch theorem over κ with its graph-theoretic and tropical analogues proved in [AC, BN, GK, MZ]. For a curve X defined over a non-Archimedean field $\mathbb{K}$, together with a semistable model $\mathfrak{X}$ for X over the valuation ring R of $\mathbb{K}$, we define a corresponding metrized complex $\mathfrak{C}\mathfrak{X}$ of curves over the residue field κ of $\mathbb{K}$ and a specialization map $\tau_*^{\mathfrak{C}\mathfrak{X}}$ from divisors on X to divisors on $\mathfrak{C}\mathfrak{X}$ preserving degrees and linear equivalence. We then establish generalizations of the specialization lemma from [B] and its weighted graph analogue from [AC], showing that the rank of a divisor cannot go down under specialization from X to $\mathfrak{C}\mathfrak{X}$. As an application, we establish a concrete link between specialization of divisors from curves to metrized complexes and the Eisenbud-Harris theory of limit linear series. Using this link, we formulate a generalization of the notion of limit linear series to curves which are not necessarily of compact type and prove, among other things, that any degeneration of a $\mathfrak{g}^r_d$ in a regular family of semistable curves is a limit $\mathfrak{g}^r_d$ on the special fiber.

As mentioned above, the specialization theory from [B] works best in the case where the family is maximally degenerate; in other cases, the dual graph of the family forgets too much information. This loss of information was partially remedied in [AC] by looking at vertex-weighted graphs, where the weight attached to a vertex is the genus of the corresponding irreducible component of the special fiber. The theory of metrized complexes of curves offers a richer solution in which one keeps track of the (normalization of the) irreducible components C_v of the special fiber themselves, and not just their genera. From the point of view of non-Archimedean geometry, this corresponds to retracting divisors on X to the *skeleton* $\Gamma_{\mathfrak{X}}$, a metric graph canonically embedded in the Berkovich analytic space X^{an}, but also keeping track, for points retracting to a vertex of $\Gamma_{\mathfrak{X}}$, of the tangent direction in which the vertex is approached. This is quite natural from the valuation-theoretic point of view: the metric graph $\Gamma_{\mathfrak{X}}$ associated to a semistable model $\mathfrak{X}$ is a canonical subset of the Berkovich analytic space X^{an}, and the curve C_v over κ is naturally homeomorphic (for each $v \in V$) to the fiber over v of the canonical retraction map $X^{\mathrm{ad}} \to X^{\mathrm{an}}$ (which sends a continuous valuation of arbitrary rank in the Huber adic space X^{ad} to its canonical rank-1 generalization). Though we could possibly (with quite a bit of additional effort) have formulated many of our theorems and proofs without mentioning Berkovich spaces, it would have resulted in a significant loss of elegance and clarity, and in any case Berkovich's theory seems ideally suited for the point of view taken here. We do not, however, require any non-trivial facts from Huber's theory of adic spaces, so we will not mention them again even though they are certainly lurking in the background. To help the reader not familiar with Berkovich's theory, we have included an expository section (Section 4) which summarizes, more or less from first principles, all of the required background facts.

There are a number of connections between the ideas in the present paper and tropical geometry; these will be explored in more detail in future work. For example, the theory of morphisms between metrized complexes of curves being developed in [ABBR] sheds interesting new light on the question of which morphisms between tropical curves are liftable. The theory developed in [BPR] shows that there is a close connection between metrized complexes of curves and "exploded tropicalizations" in the sense of [Pay, Definition 2.9]; we plan to say more about this in the future as well. As noted by Payne in [Pay], exploded tropicalizations can be thought of as an algebraic analogue of the "exploded torus fibrations" studied by Parker from a symplectic viewpoint [Par]. It could be interesting to explore connections between *loc. cit.* and the present work.

The ideas in the present paper also have Diophantine applications to the study of rational points on curves over number fields (specifically, to the method of Coleman-Chabauty). Indeed, Eric Katz and David Zureick-Brown have recently (and independently) proved a result similar to Theorem 3.4 below (Clifford's theorem for metrized complexes), and they use this result to settle a question of M. Stoll; a special case of the main result in [KZB] is the following. Let X be a smooth projective geometrically irreducible curve of genus $g \geq 2$ over $\mathbf{Q}$ and assume that the Mordell-Weil rank r of the Jacobian of X is less than g. Fix a prime number $p > 2r + 2$ and let $\mathfrak{X}$ be a proper (not necessarily semistable) regular model for X over $\mathbf{Z}_p$. Then (letting $\bar{\mathfrak{X}}^{\mathrm{sm}}$ denote the smooth locus of $\bar{\mathfrak{X}}$)

$$\#X(\mathbf{Q}) \leq \bar{\mathfrak{X}}^{\mathrm{sm}}(\mathbf{F}_p) + 2r.$$

We discuss the theorem of Katz and Zureick-Brown in more detail in Section 6.5 below, explaining how limit linear series on metrized complexes of curves can be used to illuminate their proof and put it into a broader context.

1.1. **Overview.**

We now discuss the contents of this paper in more detail.

An *edge-weighted graph* G is a connected multigraph, possibly with loop edges, having vertex set V and edge set E, and endowed with a weight (or length) function $\ell : E \to \mathbb{R}_{>0}$. A *metric graph* Γ is the geometric realization of an edge-weighted graph G in which each edge e of G is identified with a line segment of length $\ell(e)$. We call G a *model* for Γ. Subdividing an edge of G in a length-preserving fashion changes the model but not the underlying metric graph.

Let κ be an algebraically closed field. A *metrized complex* $\mathfrak{C}$ *of* κ*-curves* is the following data:

- A connected finite graph G with vertex set V and edge set E.
- A metric graph Γ having G as a model (or equivalently, a length function $\ell : E \to \mathbb{R}_{>0}$).
- For each vertex v of G, a complete, nonsingular, irreducible curve C_v of genus $g_v \geq 0$ over κ.
- For each vertex v of G, a bijection $e \mapsto x_v^e$ between the edges of G incident to v (with loop edges counted twice) and a subset $\mathcal{A}_v = \{x_v^e\}_{e \ni v}$ of $C_v(\kappa)$.

Given a metrized complex $\mathfrak{C}$ of κ-curves, there is an associated semistable curve X_0 over κ obtained by gluing the curves C_v along the points x_v^e (one intersection for each edge e of G) and forgetting the metric structure on Γ. Conversely, given a semistable curve X_0 over κ together with a positive real number for each node (which we call a "length function"), one obtains an associated metrized complex of κ-curves by letting G be the dual graph of X_0, Γ the metric graph associated to G and the given length function, C_v the normalization of the irreducible component X_v of X_0 corresponding to v, and $\mathcal{A}_v$ the preimage in C_v of the set of nodes of X_0 belonging to X_v.

Let $\mathbb{K}$ be a complete and algebraically closed non-Archimedean field with residue field isomorphic to κ and let $X/\mathbb{K}$ be a smooth, proper, connected algebraic curve. There is a metrized complex $\mathfrak{C} = \mathfrak{C}\mathfrak{X}$ canonically associated to any semistable model $\mathfrak{X}$ of X over the valuation ring R of $\mathbb{K}$: the special fiber $\bar{\mathfrak{X}}$ of $\mathfrak{X}$ is a semistable curve over κ, and one defines the length of an edge e of the dual graph of $\bar{\mathfrak{X}}$ to be the *modulus* of the open annulus $\mathrm{red}^{-1}(x^e)$, where x^e is the singular point of $\bar{\mathfrak{X}}$ corresponding to e and $\mathrm{red} : X(\mathbb{K}) \to \bar{\mathfrak{X}}(\kappa)$ is the canonical reduction map. The corresponding metric graph $\Gamma = \Gamma_{\mathfrak{X}}$ is called the *skeleton* of $\mathfrak{X}$, and there is a canonical retraction map $\tau = \tau_{\mathfrak{X}} : X^{\mathrm{an}} \twoheadrightarrow \Gamma_{\mathfrak{X}}$, where X^{an} denotes the Berkovich analytification of X. By restricting to $X(\mathbb{K})$ and then extending by linearity, we obtain a specialization map $\tau_* : \mathrm{Div}(X) \to \mathrm{Div}(\Gamma_{\mathfrak{X}})$.

If R is a discrete valuation ring and the fibered surface $\mathfrak{X}/R$ is regular, then one can use intersection-theoretic methods to define and study specialization of divisors from curves to graphs. For example, in [B], using intersection theory, a specialization homomorphism $\rho : \mathrm{Div}(X) \to \mathrm{Div}(G)$ was defined, where G is the dual graph of $\bar{\mathfrak{X}}$, with the property that $\rho(D)$ and $\tau_*(D)$ are linearly equivalent as divisors on $\Gamma_{\mathfrak{X}}$ for all $D \in \mathrm{Div}(X)$. In the present paper, we focus exclusively on τ_* and its generalization to metrized complexes of curves, though in principle all of our results can be translated (when R is a dvr and $\mathfrak{X}$ is regular) into results about this alternative specialization map.

A *divisor* on a metrized complex of curves $\mathfrak{C}$ is an element $\mathcal{D} = D_\Gamma \oplus \sum_v D_v$ of $\mathrm{Div}(\Gamma) \oplus (\oplus_v \mathrm{Div}(C_v))$ such that $\deg(D_v) = D_\Gamma(v)$ (the coefficient of v in D_Γ) for all v in V. The *degree* of $\mathcal{D}$ is defined to be the degree of D_Γ. The specialization map τ_* can be enhanced in a canonical way to a map from divisors on X to divisors on $\mathfrak{C}\mathfrak{X}$. The fiber $\tau^{-1}(v)$ of the retraction map $\tau : X(\mathbb{K}) \to \Gamma$ over a vertex $v \in V$ can be canonically identified with $\mathrm{red}^{-1}(C_v^{\mathrm{sm}}(\kappa))$, the fiber of the reduction map $\mathrm{red} : X(\mathbb{K}) \to \bar{\mathfrak{X}}(\kappa)$ over the smooth locus $C_v(\kappa) \setminus \mathcal{A}_v$ of C_v. For $P \in \tau^{-1}(v)$, we define $\bar{P}$ to be the point of $C_v(\kappa) \setminus \mathcal{A}_v$ corresponding to $\mathrm{red}(P) \in C_v^{\mathrm{sm}}(\kappa)$. With this notation, we define $\tau_*^{\mathfrak{C}\mathfrak{X}} : \mathrm{Div}(X) \to \mathrm{Div}(\mathfrak{C}\mathfrak{X})$ by setting

$$\tau_*^{\mathfrak{C}\mathfrak{X}}(P) = \begin{cases} \tau(P) & \tau(P) \notin V \\ \tau(P) \oplus \bar{P} & \tau(P) \in V \end{cases}$$

for $P \in X(\mathbb{K})$ and extending by linearity to all of $\mathrm{Div}(X)$.

Intuitively speaking, for points P of $X(\mathbb{K})$ which reduce to smooth points of $\bar{\mathfrak{X}}$, the map $\tau^{\mathfrak{C}\mathfrak{X}}$ keeps track of the reduced point $\bar{P} \in \bar{\mathfrak{X}}^{\mathrm{sm}} = \coprod (C_v(\kappa) \setminus \mathcal{A}_v)$, while if P reduces to a singular point of $\bar{\mathfrak{X}}$ then $\tau^{\mathfrak{C}\mathfrak{X}}$ instead keeps track of the retraction of P to the skeleton of the open annulus $\mathrm{red}^{-1}(P)$, which is canonically identified with the relative interior of the corresponding edge of $\Gamma_{\mathfrak{X}}$.

A *nonzero rational function* $\mathfrak{f}$ on a metrized complex of curves $\mathfrak{C}$ is the data of a rational function f_Γ on Γ and nonzero rational functions f_v on C_v for each $v \in V$. (We do *not* impose any compatibility conditions on the rational functions f_Γ and f_v.) The *divisor* of a nonzero rational function $\mathfrak{f}$ on $\mathfrak{C}$ is defined to be

$$\mathrm{div}(\mathfrak{f}) = \mathrm{div}(f_\Gamma) \oplus \sum_v \big[\mathrm{div}(f_v) + \mathrm{div}_v(f_\Gamma)\big],$$

where

$$\mathrm{div}(f_\Gamma) = \big(\sum_{u \in \Gamma} \mathrm{ord}_u(f_\Gamma)\,(u)\big), \quad \mathrm{div}_v(f_\Gamma) = \sum_{e \ni v} \mathrm{slp}_e(f_\Gamma)(x_v^e),$$

$\mathrm{slp}_e(f_\Gamma)$ denotes the outgoing slope of f_Γ at v in the direction of e, and

$$\mathrm{ord}_v(f_\Gamma) = \sum_{e \ni v} \mathrm{slp}_e(f_\Gamma).$$

Divisors of the form $\mathrm{div}(\mathfrak{f})$ are called *principal*, and the principal divisors form a subgroup of $\mathrm{Div}^0(\mathfrak{C})$, the group of divisors of degree zero on $\mathfrak{C}$. Two divisors in $\mathrm{Div}(\mathfrak{C})$ are called *linearly equivalent* if they differ by a principal divisor.

The motivation for our definitions of $\tau_*^{\mathfrak{C}\mathfrak{X}}$ and $\mathrm{div}(\mathfrak{f})$ come from the following fundamental relation, a consequence of the non-Archimedean Poincaré-Lelong formula due to Thuillier (see §5 of [BPR]). Let f be a nonzero rational function on X and let $\mathfrak{f}$ be the corresponding nonzero rational function on $\mathfrak{C}\mathfrak{X}$, where f_Γ is the restriction to Γ of the piecewise linear function $\log|f|$ on X^{an} and $f_v \in \kappa(C_v)$ for $v \in V$ is the reduction of f_v to C_v. Then

$$\tau_*^{\mathfrak{C}\mathfrak{X}}(\mathrm{div}(f)) = \mathrm{div}(\mathfrak{f}).$$

In particular, we have $\tau_*^{\mathfrak{C}\mathfrak{X}}(\mathrm{Prin}(X)) \subseteq \mathrm{Prin}(\mathfrak{C}\mathfrak{X})$.

A divisor $\mathcal{E} = E_\Gamma \oplus \sum_v E_v$ is called *effective* if E_Γ and the divisors E_v are all effective. The *rank* $r_{\mathfrak{C}}$ of a divisor $\mathcal{D} = D_\Gamma \oplus \sum_v D_v$ in $\mathrm{Div}(\mathfrak{C})$ is defined to be the largest integer k

such that $\mathcal{D} - \mathcal{E}$ is linearly equivalent to an effective divisor for all effective divisors $\mathcal{E}$ of degree k on $\mathfrak{C}$ (so in particular $r_{\mathfrak{C}}(\mathcal{D}) \geq 0$ if $\mathcal{D}$ is linearly equivalent to an effective divisor, and otherwise $r_{\mathfrak{C}}(\mathcal{D}) = -1$).

If $r_X(D)$ denotes the usual rank function $r_X(D) = \dim |D| = h^0(D) - 1$ on $\mathrm{Div}(X)$, there is an important semicontinuity result relating r_X to $r_{\mathfrak{C}\mathfrak{X}}$:

Theorem 1.1 (Specialization Theorem)**.** *For all $D \in \mathrm{Div}(X)$, we have*

$$r_X(D) \leq r_{\mathfrak{C}\mathfrak{X}}(\tau_*^{\mathfrak{C}\mathfrak{X}}(D)).$$

Since $r_{\mathfrak{C}\mathfrak{X}}(\tau_*^{\mathfrak{C}\mathfrak{X}}(D)) \leq r_\Gamma(\tau_*(D_\Gamma))$, Theorem 1.1 is a strengthening of the analogous specialization result from [B].

The theory of divisors and linear equivalence on metrized complexes of curves generalizes both the classical theory for algebraic curves and the corresponding theory for metric graphs and tropical curves found in [GK, MZ]. The former corresponds to the case where G consists of a single vertex v and no edges and $C = C_v$ is an arbitrary smooth curve. The latter corresponds to the case where the curves C_v have genus zero for all $v \in V$. Since any two points on a curve of genus zero are linearly equivalent, it is easy to see that the divisor theory (and rank function) on $\mathfrak{C}$ and Γ are essentially equivalent. More precisely, two divisors $\mathcal{D} = D_\Gamma \oplus \sum_v D_v$ and $\mathcal{D}' = D'_\Gamma \oplus \sum_v D'_v$ on $\mathfrak{C}$ are linearly equivalent if and only if D_Γ and D'_Γ are linearly equivalent on Γ, and for every $\mathcal{D} \in \mathrm{Div}(\mathfrak{C})$ we have $r_{\mathfrak{C}}(\mathcal{D}) = r_\Gamma(D_\Gamma)$. Obviously, in the presence of higher genus curves among the C_v, the divisor theory on Γ and $\mathfrak{C}$ can be very different. In addition, different choices of $\mathcal{A}_v$ can drastically change both the linear equivalence relation and the rank function.

Theorem 1.1 refines, and provides a new proof for, the specialization lemma for vertex-weighted graphs from [AC]. Our metric graph generalization of the latter result is as follows. A (*vertex-*)*weighted metric graph* is a metric graph Γ together with a function $\omega : \Gamma \to \mathbb{Z}_{\geq 0}$ such that $\omega(x) = 0$ for all but finitely many $x \in \Gamma$. The *genus* of (Γ, ω) is defined to be $g^\# = g(\Gamma) + \sum_{x \in \Gamma} \omega(x)$. Note that if $\Gamma^\#$ is any metric graph obtained from Γ by attaching $\omega(x)$ loops of arbitrary positive length at each point x where $\omega(x) > 0$, then $g^\#$ is equal to the genus of $\Gamma^\#$. If $D \in \mathrm{Div}(\Gamma)$, we denote by $r^\#(D)$ the rank of D considered (in the obvious way) as a divisor on the metric graph $\Gamma^\#$. If we define $\mathcal{W}$ to be the divisor $\sum_{x \in \Gamma} \omega(x)(x)$ on Γ, then

$$r^\#(D) = \min_{0 \leq E \leq \mathcal{W}} \left(\deg(E) + r_\Gamma(D - 2E)\right).$$

In particular, $r^\#(D)$ is independent of the choice of $\Gamma^\#$ and there is a formula for calculating it which is intrinsic to Γ. With this terminology, the specialization lemma for weighted metric graphs can be stated as follows:

Theorem 1.2. *Let $\mathfrak{X}$ be a semistable model for X and let (Γ, ω) be the corresponding weighted metric graph, where $\Gamma = \Gamma_{\mathfrak{X}}$ and $\omega(x) = g_x$ is the genus of C_v if $x = v \in V$ and is zero otherwise. Then for every $D \in \mathrm{Div}(X)$,*

$$r_X(D) \leq r^\#(\tau_*(D)).$$

Using the theory of rank-determining sets, as discussed in Appendix A, the specialization lemma for (non-metric) weighted graphs proved in [AC] is a direct consequence of Theorem 1.2.

A useful refinement of Theorem 1.2 is as follows. Given a metrized complex of κ-curves $\mathfrak{C}$ and a divisor $\mathcal{D} = D_\Gamma \oplus \sum D_v$ in $\mathrm{Div}(\mathfrak{C})$, we define for each $v \in V$ a function $\eta_v : \mathbb{Z}_{\geq 0} \to \mathbb{Z}_{\geq 0}$ by letting $\eta_v(k)$ be the smallest nonnegative integer n such that there exists a divisor $\widetilde{D}_v$ supported on $\mathcal{A}_v$ with $r_{C_v}(D_v + \widetilde{D}_v) = k$ and $\deg(D_v + \widetilde{D}_v) = n$. For an effective divisor $E \in \mathrm{Div}(\Gamma)$ supported on V, define

$$\eta(E) = \sum_{v \in V} \eta_v(E(v))(v).$$

(Note that the definition of η depends on the collection $\{D_v\}$ but not on D_Γ.)

Theorem 1.3 (Refined weighted specialization inequality)**.** *With hypotheses as in Theorem 1.2, for every divisor $D \in \mathrm{Div}(X)$ we have*

$$r_X(D) \leq r_{\mathfrak{CX}}(\tau_*^{\mathfrak{CX}}(D)) \leq \min_{E \geq 0} \left(\deg(E) + r_\Gamma(\tau_*(D) - \eta(E))\right).$$

By Clifford's theorem applied to the curve C_v, we have $\eta_v(k) \geq 2k$ for $0 \leq k \leq g_v$ and $\eta_v(k) \geq k + g_v$ for $k \geq g_v$. Using this observation, Theorem 1.2 can be deduced easily from Theorem 1.3. An advantage of Theorem 1.3 is that it gives a much stronger result in the "generic" case. More precisely, one can deduce from Theorem 1.3 the following:

Corollary 1.4. *If the marked curves $(C_v, \mathcal{A}_v)$ appearing in the data for $\mathfrak{CX}$ are general in their moduli, then for a general divisor $D \in \mathrm{Div}(X)$ with $\tau_*(D) \geq \mathcal{W}$, we have*

$$r_X(D) \leq r_\Gamma(\tau_*(D) - \mathcal{W}).$$

As with the corresponding specialization theory from [B], the main utility of Theorem 1.1 is that the rank function $r_{\mathfrak{C}}$ on a metrized complex of curves is surprisingly well-behaved; for example, it satisfies an analogue of the Riemann-Roch formula. In order to state this result, we need to define the genus and canonical class. The *genus* of a metrized complex of curves $\mathfrak{C}$, denoted $g(\mathfrak{C})$, is by definition $g(\mathfrak{C}) = g(\Gamma) + \sum_{v \in V} g_v$, where g_v is the genus of C_v. A *canonical divisor* on $\mathfrak{C}$, denoted $\mathcal{K}$, is defined to be any divisor linearly equivalent to $K^{\#} \oplus \sum_{v \in V}(K_v + A_v)$, where $K^{\#} = \sum_v \left(\deg_G(v) + 2g_v - 2\right)(v)$, K_v is a canonical divisor on C_v, and A_v is the sum of the $\deg_G(v)$ points in $\mathcal{A}_v$. The following result generalizes both the classical Riemann-Roch theorem for algebraic curves and the Riemann-Roch theorem for metric graphs:

Theorem 1.5 (Riemann-Roch for metrized complexes of algebraic curves)**.** *Let $\mathfrak{C}$ be a metrized complex of algebraic curves over κ and $\mathcal{K}$ a divisor in the canonical class of $\mathfrak{C}$. For any divisor $\mathcal{D} \in \mathrm{Div}(\mathfrak{C})$, we have*

$$r_{\mathfrak{C}}(\mathcal{D}) - r_{\mathfrak{C}}(\mathcal{K} - \mathcal{D}) = \deg(\mathcal{D}) - g(\mathfrak{C}) + 1.$$

The proof of Theorem 1.5 makes use of a suitable notion of *reduced divisors* for metrized complexes of curves in order to verify the abstract Riemann-Roch criterion from [BN]. A sketch of the proof is as follows.

Fix a point $v_0 \in \Gamma$ (which might or might not be in V) and let $\Gamma' = \Gamma \setminus \{v_0\}$. For a closed connected set $S \subset \Gamma$ with topological boundary ∂S and $x \in \partial S$ a boundary point, we denote by $\mathrm{outdeg}_S(x)$ the number of "outgoing" tangent directions at x (i.e., tangent directions emanating outward from S). Moreover, for $v \in V \cap \partial S$ we denote by $\mathrm{div}_v(\partial S)$ the divisor $\sum_{e \in \mathrm{out}_S(v)}(x_v^e)$ on C_v, where $\mathrm{out}_S(v)$ is the collection of edges e incident to v which correspond to outgoing tangent directions. A divisor $\mathcal{D} = D_\Gamma + \sum_{v \in V} D_v$ on $\mathfrak{C}$ is called v_0*-reduced* if each of the following three conditions hold:

- $D_\Gamma(x) \geq 0$ for all $x \in \Gamma'$.
- D_v is linearly equivalent on C_v to an effective divisor for all $v \in V \cap \Gamma'$.
- For every closed connected set $S \subset \Gamma'$, there exists $x \in \partial S$ such that either
  - (a) $x \notin V$ and $D_\Gamma(x) - \mathrm{outdeg}_S(x) < 0$, or
  - (b) $x = v \in V$ and $D_v - \mathrm{div}_v(\partial S)$ is not equivalent to an effective divisor on C_v.

By analogy with previous results concerning divisors on graphs and metric graphs (see e.g. [A, BN, HKN]), we establish the following existence and uniqueness result for reduced divisors on metrized complexes of curves:

Theorem 1.6. *Every divisor $\mathcal{D} \in \mathrm{Div}(\mathfrak{C})$ is linearly equivalent on $\mathfrak{C}$ to a v_0-reduced divisor. Moreover, if $\mathcal{D}' = D'_\Gamma \oplus \sum_{v \in V} D'_v$ and $\mathcal{D}'' = D''_\Gamma \oplus \sum_{v \in V} D''_v$ are v_0-reduced divisors which are both linearly equivalent to $\mathcal{D}$, then $D'_\Gamma = D''_\Gamma$ and D'_v is linearly equivalent to D''_v for all $v \in V$.*

An *acyclic orientation* on Γ is an acyclic orientation on some model of Γ, i.e., a decomposition of Γ into closed directed edges with disjoint interiors such that no directed cycles are formed. Given an acyclic orientation π on Γ, we denote by $\deg^+_\pi(x)$ the number of tangent directions emanating from x which are compatible with π. Moreover, for $v \in V$ we denote by $E^+(v)$ the set of edges incident to v which (locally near v) are oriented outward from v. A *moderator* on $\mathfrak{C}$ is a divisor $\mathcal{M} \in \mathrm{Div}(\mathfrak{C})$ of the form $\mathcal{D}^{\pi,\{D_v\}} = D^\pi_\Gamma \oplus \sum D^\pi_v$, where π is an acyclic orientation of Γ and for each $v \in V$, D_v is a divisor of degree $g_v - 1$ on C_v which is not equivalent to an effective divisor. Here we set

$$D^\pi_\Gamma = \sum_{x \in \Gamma} \left(\deg^+_\pi(x) - g_x - 1\right)(x),$$

where g_v is the genus of C_v if $v \in V$ and is zero otherwise, and

$$D^\pi_v = D_v + \sum_{e \in E^+(v)} (x^e_v).$$

The *dual moderator* $\bar{\mathcal{M}}$ is defined to be $\mathcal{D}^{\bar{\pi},\{K_v - D_v\}}$, where $\bar{\pi}$ is obtained from π by reversing the orientation of every oriented segment. It is easy to see that $\deg(\mathcal{M}) = g(\mathfrak{C}) - 1$ for any moderator $\mathcal{M}$, and that $\mathcal{M} + \bar{\mathcal{M}}$ belongs to the canonical class on $\mathfrak{C}$. Moreover, we show that a moderator can never be equivalent to an effective divisor.

The main application of the existence and quasi-uniqueness of v_0-reduced divisors on $\mathfrak{C}$ is the following result which (using the abstract Riemann-Roch criterion from [BN]) directly implies Theorem 1.5 :

Theorem 1.7. *For all $\mathcal{D} \in \mathrm{Div}(\mathfrak{C})$, exactly one of the following holds:*

- *$\mathcal{D}$ is equivalent to an effective divisor.*
- *$\mathcal{M} - \mathcal{D}$ is equivalent to an effective divisor for some moderator $\mathcal{M}$.*

Taken together, Theorems 1.1 and 1.5 have some interesting consequences. For example, the specialization theorem implies that the specialization to $\mathfrak{CX}$ of any canonical divisor on X has degree $2g - 2$ and rank at least $g - 1$ (where $g = g(X) = g(\mathfrak{CX})$), while the Riemann-Roch theorem easily implies that any such divisor on $\mathfrak{CX}$ belongs to the canonical divisor class. Thus we have:

Theorem 1.8. *For any canonical divisor K_X on X, its specialization $\tau^{\mathfrak{CX}}_*(K_X)$ belongs to the canonical divisor class on $\mathfrak{CX}$.*

As a corollary, Theorem 1.8 implies that $\tau_*(K_X)$ is a canonical divisor on Γ. This corollary was proved in [B], for discretely valued R, by a completely different argument based on the adjunction formula for arithmetic surfaces. It does not seem easy to prove the general (not necessarily discretely valued) case using intersection theory.

Our theory of linear series on metrized complexes of curves has close connections with the Eisenbud-Harris theory of limit linear series for strongly semistable curves of compact type, and suggests a way to generalize the Eisenbud-Harris theory to more general semistable curves. In order to state the connection, we first need to generalize the definition of the rank function $r_{\mathfrak{C}}$. Let $\mathcal{F} = \{F_v \ : \ v \in V\}$, where F_v is a κ-subspace of the function field $\kappa(C_v)$ for each $v \in V$. For $\mathcal{D} \in \operatorname{Div}(\mathfrak{C})$, we define the $\mathcal{F}$-rank of $\mathcal{D}$, denoted $r_{\mathfrak{C},\mathcal{F}}(D)$, to be the largest integer k such that for any effective divisor $\mathcal{E}$ of degree k on $\mathfrak{C}$, there is a rational function $\mathfrak{f} = (f_\Gamma, (f_v)_{v\in V})$ with $f_v \in F_v$ for all $v \in V$ such that $\mathcal{D} - \mathcal{E} + \operatorname{div}(\mathfrak{f}) \geq 0$.

A proper nodal curve X_0 over κ is of *compact type* if its dual graph G is a tree. For such curves, Eisenbud and Harris define a notion of (crude) *limit* $\mathfrak{g}^r_d$, which we discuss in detail in Section 6.2 below. A *crude limit* $\mathfrak{g}^r_d$ L on X_0 is the data of a (not necessarily complete) degree d and rank r linear series L_v on X_v for each vertex $v \in V$ such that if two components X_u and X_v of X_0 meet at a node p, then for any $0 \leq i \leq r$,

$$a_i^{L_v}(p) + a_{r-i}^{L_u}(p) \ \geq \ d\,,$$

where $a_i^L(p)$ denotes the i^{th} term in the vanishing sequence of a linear series L at p. A crude limit series is *refined* if all the inequalities in the above definition are equalities. For simplicity, all limit linear series in the remainder of this introduction will be crude.

We can canonically associate to a proper strongly semistable curve X_0 a metrized complex $\mathfrak{C}X_0$ of κ-curves, called the *regularization* of X_0, by assigning a length of 1 to each edge of G, and we write X_v for the irreducible component of X_0 corresponding to a vertex $v \in V$. (This is the metrized complex associated to any regular smoothing $\mathfrak{X}$ of X_0 over any discrete valuation ring R with residue field κ.)

Theorem 1.9. *Let $\mathfrak{C}X_0$ be the metrized complex of curves associated to a proper strongly semistable curve X_0/κ of compact type. Then there is a bijective correspondence between the following:*

- *Crude limit $\mathfrak{g}^r_d$'s on X_0.*
- *Equivalence classes of pairs $(\mathcal{H}, \mathcal{D})$, where $\mathcal{H} = \{H_v\}$, H_v is an $(r+1)$-dimensional subspace of $\kappa(X_v)$ for each $v \in V$, and $\mathcal{D} \in \operatorname{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$ is a divisor of degree d on $\mathfrak{C}X_0$ with $r_{\mathfrak{C}X_0,\mathcal{H}}(\mathcal{D}) = r$. Here we say that $(\mathcal{H}, \mathcal{D}) \sim (\mathcal{H}', \mathcal{D}')$ if there is a rational function $\mathfrak{f} = (f_\Gamma, (f_v)_{v\in V})$ on $\mathfrak{C}X_0$ such that $D' = D + \operatorname{div}(\mathfrak{f})$ and $H_v = H'_v \cdot f_v$ for all $v \in V$.*

Theorem 1.9, combined with our Riemann-Roch theorem for metrized complexes of curves, provides an arguably more conceptual proof of the fact (originally established in [EH]) that limit linear series satisfy analogues of the classical theorems of Riemann and Clifford. The point is that $r_{\mathfrak{C}X_0,\mathcal{H}}(\mathcal{D}) \leq r_{\mathfrak{C}X_0}(\mathcal{D})$ for all $\mathcal{D} \in \operatorname{Div}(\mathfrak{C}X_0)$ and therefore upper bounds on $r_{\mathfrak{C}X_0}(\mathcal{D})$ which follow from Riemann-Roch imply corresponding upper bounds on the restricted rank $r_{\mathfrak{C}X_0,\mathcal{H}}(\mathcal{D})$.

Motivated by Theorem 1.9, we propose the following definition. Let X_0 be a proper strongly semistable (but not necessarily compact type) curve over κ with regularization $\mathfrak{C}X_0$. A *limit* $\mathfrak{g}^r_d$ on X_0 is an equivalence class of pairs $(\mathcal{H} = \{H_v\}, \mathcal{D})$ as above, where H_v

is an $(r+1)$-dimensional subspace of $\kappa(X_v)$ for each $v \in V$, and $\mathcal{D}$ is a degree d divisor on $\mathfrak{C}X_0$ with $r_{\mathfrak{C}X_0,\mathcal{H}}(\mathcal{D}) = r$.

As a partial justification for this definition, we prove that if R is a dvr with residue field κ and $\mathfrak{X}/R$ is a regular arithmetic surface whose generic fiber X is smooth and whose special fiber X_0 is strongly semistable, then for any divisor D on X with $r_X(D) = r$ and $\deg(D) = d$, our specialization machine gives rise in a natural way to a limit $\mathfrak{g}^r_d$ on X_0 (see Theorem 6.8 below for a precise statement).

In Section 7, we illustrate the theory of divisors and limit linear series on metrized complexes of curves via a detailed analysis of the case of a generic chain of genus-one curves. The description of ranks of divisors on such metrized complexes turn out to be quite similar to the example considered in Cools et. al. [CDPR] of a generic chain of genus-one metric graphs. Our analysis of this example provides some insight into the relationship between the proofs given in Eisenbud-Harris [EH] and [CDPR] of the non-existence part of the Brill-Noether theorem in algebraic geometry. In addition, we show that every divisor of degree d and rank r on such a metrized complex can be completed to a limit g^r_d, something which does not hold for metrized complexes in general.

2. Metrized complexes of algebraic curves

Let κ be an algebraically closed field of arbitrary characteristic and let $\mathfrak{C}$ be a metrized complex of κ-curves, as defined in Section 1.1.

2.1. Divisors and their rank.

Recall that the group of divisors on $\mathfrak{C}$, denoted by $\mathrm{Div}(\mathfrak{C})$, is the subgroup of the direct sum $\mathrm{Div}(\Gamma) \oplus (\oplus_{v\in V} \mathrm{Div}(C_v))$ consisting of all the elements $\mathcal{D} = D_\Gamma \oplus \sum_{v\in V} D_v$ such that $\deg(D_v) = D_\Gamma(v)$ (the coefficient of v in D_Γ). The degree map $\deg : \mathrm{Div}(\mathfrak{C}) \to \mathbb{Z}$ is defined by $\deg(D_\Gamma \oplus \sum_v D_v) = \deg(D_\Gamma)$.

Remark 2.1. While it is true that $\mathrm{Div}(\mathfrak{C})$ is canonically isomorphic to the free abelian group generated by the union of the points of $\Gamma \setminus V$ and the κ-points of C_v, for $v \in V$, the presentation we give for $\mathrm{Div}(\mathfrak{C})$ makes it easier to break the study of divisors on $\mathfrak{C}$ into a combinatorial part corresponding to Γ and an algebraic part corresponding to the curves C_v.

If $\mathcal{G}$ is a subgroup of $\mathbb{R}$ containing the lengths of all edges of G, we write $\mathrm{Div}(\mathfrak{C})_\mathcal{G}$ for the subgroup of $\mathrm{Div}(\mathfrak{C})$ consisting of all divisors $\mathcal{D} = D_\Gamma \oplus \sum_v D_v$ such that $\mathrm{supp}(D_\Gamma)$ consists entirely of $\mathcal{G}$-*rational points* of Γ (i.e., points of Γ whose distances to the vertices of G are in $\mathcal{G}$).

Recall from Section 1.1 that the *rank* $r_\mathfrak{C}$ of a divisor $\mathcal{D} = D_\Gamma + \sum_v D_v$ on $\mathfrak{C}$ is defined to be the largest integer k such that $\mathcal{D} - \mathcal{E}$ is linearly equivalent to an effective divisor for all effective divisors $\mathcal{E}$ of degree k on $\mathfrak{C}$. In particular, $r_\mathfrak{C}(\mathcal{D}) \geq 0$ iff $\mathcal{D}$ is linearly equivalent to an effective divisor, and otherwise $r_\mathfrak{C}(\mathcal{D}) = -1$.

If $\mathcal{G}$ is any subgroup of $\mathbb{R}$ containing the lengths of all edges of G, one can define an analogous rank function $r_{\mathfrak{C},\mathcal{G}}$ for divisors in $\mathrm{Div}(\mathfrak{C})_\mathcal{G}$ by restricting the effective divisors $\mathcal{E}$ in the above definition to lie in $\mathrm{Div}(\mathfrak{C})_\mathcal{G}$. As a consequence of the results in Appendix A, we will see that $r_{\mathfrak{C},\mathcal{G}}(\mathcal{D}) = r_\mathfrak{C}(\mathcal{D})$ for any divisor $\mathcal{D} \in \mathrm{Div}(\mathfrak{C})_\mathcal{G}$, so the restriction in the definition of $r_{\mathfrak{C},\mathcal{G}}$ does not affect the rank of divisors.

Remark 2.2. Divisor theory on metrized complexes generalizes both divisor theory on metric graphs and divisor theory on smooth curves. The latter corresponds to the case where G consists of a single vertex v and no edges and $C = C_v$ is an arbitrary proper smooth curve over κ. The former corresponds to the case where $g_v = 0$ for all $v \in V$. Since in this case any two divisors of the same degree on each C_v are linearly equivalent, it is easy to see that the notion of linear equivalence (and rank) on $\mathfrak{C}$ and Γ are essentially equivalent. More precisely, two divisors $\mathcal{D} = D_\Gamma \oplus \sum_v D_v$ and $\mathcal{D}' = D'_\Gamma \oplus \sum_v D'_v$ on $\mathfrak{C}$ are linearly equivalent if and only if D_Γ and D'_Γ are linearly equivalent on Γ. In addition, for every $\mathcal{D} \in \mathrm{Div}(\mathfrak{C})$ we have $r_{\mathfrak{C}}(\mathcal{D}) = r_\Gamma(D_\Gamma)$.

2.2. **Regularization of nodal curves.** Let X_0 be a (reduced) connected projective nodal curve over κ with irreducible components $X_{v_1}, \dots, X_{v_k}$. We can associate to X_0 a metrized complex $\mathfrak{C}X_0$, called the *regularization* of X_0, as follows. The underlying metric graph Γ_0 has model the dual graph $G_0 = (V_0, E_0)$ of X_0. Recall that the vertex set V_0 of G_0 consists of vertices $v_1, \dots, v_k$ (in bijection with the irreducible components of X_0) and the edges of G_0 are in bijection with singular points of X_0. (Note that G_0 might have loop edges.) The length of all the edges in Γ_0 are equal to one. The κ-curve C_{v_i} in $\mathfrak{C}X_0$ is the normalization of X_i for $i = 1, \dots, k$. The collection $\mathcal{A}_{v_i}$ is the set of all κ-points of C_{v_i} which lie over a singular point of X_0 in X_{v_i}. By the definition of the dual graph, these points are in bijection with the edges adjacent to v_i in G_0.

Suppose now that X_0 is strongly semistable (so that G_0 does not have any loop edge and $C_{v_i} = X_{v_i}$). The rank function $r_{\mathfrak{C},\mathbb{Z}}$ on $\mathrm{Div}(\mathfrak{C}X_0)_\mathbb{Z}$ can be reformulated as follows. Let $\mathrm{Pic}(X_0)$ be the Picard group of X_0, and consider the restriction map $\pi : \mathrm{Pic}(X_0) \to \oplus_i \mathrm{Pic}(X_{v_i})$. For any line bundle $\mathcal{L}$ on X_0, and $v \in V_0$, denote by $\mathcal{L}_v$ the restriction of $\mathcal{L}$ to X_v.

Two line bundles $\mathcal{L}$ and $\mathcal{L}'$ in $\mathrm{Pic}(X_0)$ are said to be *combinatorially equivalent* if there exists a function $f : V_0 \to \mathbb{Z}$ such that $\mathcal{L}'_v = \mathcal{L}_v(\mathrm{div}_v(f))$ in $\mathrm{Pic}(X_v)$ for any vertex v of G_0. (In particular, two line bundles $\mathcal{L}$ and $\mathcal{L}'$ in $\mathrm{Pic}(X_0)$ with $\pi(\mathcal{L}) = \pi(\mathcal{L}')$ are combinatorially equivalent.) Recall that, as before, the definition of div_v is as follows:

$$\mathrm{div}_v(f) = \sum_{u:\{u,v\}\in E_0} (f(u) - f(v))(x_v^{\{u,v\}})$$

where $x_v^{\{u,v\}}$ is the point of C_v labelled with the edge $\{u, v\}$. Note that the function $f : V_0 \to \mathbb{Z}$ in the definition of combinatorially equivalent line bundles above is unique up to an additive constant. For such an f, we denote by $\mathcal{L}^f$ any line bundle $\mathcal{L}'$ in $\mathrm{Pic}(X_0)$ such that $\mathcal{L}'_v = \mathcal{L}_v(\mathrm{div}_v(f))$ for all $v \in V_0$.

For a line bundle $\mathcal{L}$ on X_0, define the *combinatorial rank* of $\mathcal{L}$, denoted $r_c(\mathcal{L})$, to be the maximum integer r such that any effective divisor $E = \sum_{v\in V_0} E(v)(v)$ on G_0 of degree r, there is a line bundle $\mathcal{L}'$ combinatorially equivalent to $\mathcal{L}$ with

$$\dim_\kappa H^0(X_v, \mathcal{L}'_v) \geq E(v) + 1$$

for all $v \in V$. In particular, if no non-negative integer r with the above property exists, the combinatorial rank of $\mathcal{L}$ will be equal to -1. Obviously, the combinatorial rank of $\mathcal{L}$ only depends on the combinatorially equivalence class of $\mathcal{L}$.

Let $\mathcal{D} = D_\Gamma \oplus \Sigma_v D_v$ be a divisor in $\mathrm{Div}(\mathfrak{C}X_0)_\mathbb{Z}$. By $\mathcal{L}(\mathcal{D})$ we denote any line bundle on X_0 which has components $\mathcal{L}(D_v)$ on each X_v for any vertex v of G_0.

Proposition 2.3. *Let* X_0 *be a strongly semistable curve over* κ *and let* $\mathfrak{C}X_0$ *be the corresponding metrized complex. Then:*

1. *Two divisors* $\mathcal{D}$ *and* $\mathcal{D}'$ *in* $\mathrm{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$ *are linearly equivalent iff* $\mathcal{L}(\mathcal{D})$ *and* $\mathcal{L}(\mathcal{D}')$ *are combinatorially equivalent.*
2. *For any divisor* $\mathcal{D} \in \mathrm{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$*, the combinatorial rank of* $\mathcal{L}(\mathcal{D})$ *is equal to* $r_{\mathfrak{C}X_0}(\mathcal{D})$*.*

Proof. The first part follows by definition. To see the second part, let r be the combinatorial rank of $\mathcal{L}(\mathcal{D})$. As we mentioned before, we will prove in Appendix A that $r_{\mathfrak{C}X_0}(\mathcal{D}) = r_{\mathfrak{C}X_0,\mathbb{Z}}(\mathcal{D})$. Thus, it will be enough to show that $r_{\mathfrak{C}X_0,\mathbb{Z}}(\mathcal{D}) = r$.

We first prove that $r_{CX_0,\mathbb{Z}}(\mathcal{D}) \geq r$. Let $\mathcal{E} = E \oplus \sum_v E_v$ be an effective divisor in $\mathrm{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$. By the definition of the combinatorial rank, there exists a function $f : V_0 \to \mathbb{Z}$ such that

$$\dim_\kappa H^0(X_v, \mathcal{L}_v^f) \geq E(v) + 1 \qquad \text{for any vertex } v \text{ in } G_0.$$

Since $\deg(E_v) = E(v)$, this implies the existence of a global section f_v of $\mathcal{L}_v^f$ such that f_v has order of vanishing at least $E_v(x)$ at any point x of X_v.

Let f_Γ be the rational function defined by f on Γ. (Recall that $f_\Gamma|_{V_0} = f$ and f is integer affine with slope $f(u) - f(v)$ on each edge $\{u, v\}$.) For the rational function $\mathfrak{f}$ on $\mathfrak{C}X_0$ defined by f_Γ and f_v, $v \in V_0$, one has $\mathrm{div}(\mathfrak{f}) + \mathcal{D} - \mathcal{E} \geq 0$. This shows that $r_{CX_0,\mathbb{Z}}(\mathcal{D}) \geq r$.

We now prove that $r_{CX_0,\mathbb{Z}}(\mathcal{D}) \leq r$. Let $\mathcal{L} = \mathcal{L}(D)$ be the line bundle on X_0 defined by $\mathcal{D}$. Let E be an effective divisor of degree $r_{CX_0,\mathbb{Z}}(\mathcal{D})$ on G_0 and let $\mathcal{E} = E \oplus E_v$ be an effective divisor in $\mathrm{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$. There exists a rational function $\mathfrak{f}$ with $\mathrm{div}(\mathfrak{f}) \in \mathrm{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$ such that $\mathrm{div}(\mathfrak{f}) + \mathcal{D} - \mathcal{E} \geq 0$. In particular, if f_Γ denotes the Γ-component of $\mathfrak{f}$, f_Γ is linear on each edge of G_0 and in addition $\mathrm{div}(f_\Gamma) + D \geq 0$. Up to an additive constant, there are only a finite number of f_Γ with these properties. In particular, since κ is infinite, for a generic choice of effective divisors E_v of degree $E(v)$ on X_v, one can assume that $\mathrm{div}(f_\Gamma)$ is fixed. In other words, for generic choices of E_v, all the rational functions f_v define global sections of $\mathcal{L}_v^f$, and so $\dim_\kappa H^0(X_v, \mathcal{L}_v^f) \geq \deg(E_v) + 1 = E(v) + 1$. Since this holds for any effective divisor E of degree $r_{\mathfrak{C}X_0,\mathbb{Z}}(\mathcal{D})$ on G_0, we obtain the desired inequality. □

In Section 4, we will see how to associate a metrized complex to any semistable model of a proper smooth curve X over an algebraically closed and complete non-Archimedean field $\mathbb{K}$.

2.3. **Metrized complexes over an arbitrary ground field.** For simplicity, we assume throughout the paper that κ is algebraically closed. However, let us quickly indicate how the theory of metrized complexes of curves can be developed over an arbitrary field κ.

Let κ be a field. Let (Γ, ω) be a metric graph with a model $G = (V, E)$, and C_v a collection of geometrically connected, smooth, projective κ-curves of genus $g_v = \omega(v)$ for each vertex v of G. We suppose as before that for each vertex v a collection $\mathcal{A}_v$ of $\deg_G(v)$ closed points of C_v is given, and that these points are labeled with the edges adjacent to v. For an edge $e = \{u, v\}$ in G, and for the two corresponding closed points x_u^e and x_v^e of C_u and C_v labeled with e, we assume in addition that an isomorphism ψ_e of κ-extensions $\kappa(x_u^e) \cong \kappa(x_v^e)$ is given. We let δ be the function assigning to each edge $e = \{u, v\}$ of G the integer weight $[\kappa(x_u^e) : \kappa] = [\kappa(x_v^e) : \kappa]$. A *metrized complex of* κ*-curves* consists of the data $\{\Gamma, G, \{C_v\}, \{\mathcal{A}_v\}, \{\psi_e\}, \delta\}$.

Over a general base field κ, the definition of $\mathrm{Div}(\mathfrak{C})$ remains the same. However, the definition of the principal divisor $\mathrm{div}(\mathfrak{f})$ associated to a nonzero rational function $\mathfrak{f}$ on $\mathfrak{C}$ has to be modified by taking into account the edge-weights of Γ. This simply consists in modifying the term $\mathrm{ord}_u(f_\Gamma)$ in the above definition, for the rational function f_Γ on Γ, by taking the weighted sum of the slopes of f_Γ along the tangent directions to Γ at u. The degree map $\deg : \mathrm{Div}(\mathfrak{C}) \to \mathbb{Z}$ must also be modified, taking into account the weights of edges: for a point x lying in the interior of an edge e in G, we put $\deg((x)) = \delta(e)$.

A metrized complex of curves $\mathfrak{C}$ over κ gives rise by base change from κ to κ^{sep} to a metrized complex $\mathfrak{C}^{\mathrm{sep}}$ over κ^{sep}, the separable closure of κ. The underlying metric graph Γ^{sep} of $\mathfrak{C}^{\mathrm{sep}}$ has model $G^{\mathrm{sep}} = (V, E^{\mathrm{sep}})$ obtained from the model $G = (V, E)$ of Γ in which each edge $e = \{u, v\}$ is replaced by $[\kappa(x_u^e) : \kappa] = [\kappa(x_v^e) : \kappa]$ parallel edges. The length of all these edges in Γ^{sep} are identical to the length of the original edge e in Γ. The curves C_v^{sep} are obtained from C_v by base change from κ to κ^{sep}. The collection $\mathcal{A}_v^{\mathrm{sep}}$ consists of all the geometric points in C_v^{sep} which lie over a point of $\mathcal{A}_v$. Obviously, the number of these points is $\deg_{G^{\mathrm{sep}}}(v)$. The given isomorphisms of κ-extensions $\kappa(x_u^e) \sim \kappa(x_v^e)$, for the two closed points x_u^e and x_v^e of C_u and C_v labeled with the same edge $e = \{u, v\}$ in G, allows us to provide a labeling of the points of $\mathcal{A}_v^{\mathrm{sep}}$ with th e edges of G^{sep} incident to v in a coherent way.

There is a natural injective map $\iota : \mathrm{Div}(\mathfrak{C}) \longrightarrow \mathrm{Div}(\mathfrak{C}^{sep})$ which on each curve C_v sends a closed point x of C_v to the sum of the points y in C_v^{sep} lying above x, and on Γ sends each point $P \in \Gamma$ to the sum of all the copies of P in Γ^{sep}. It is easy to see that ι respects the degree of divisors and sends principal divisors on $\mathfrak{C}$ to principal divisors on $\mathfrak{C}^{sep}$. It follows from the results of Appendix A that rank of a divisor $\mathcal{D}$ on $\mathfrak{C}$ is equal to the rank of $\iota(\mathcal{D})$ on $\mathfrak{C}^{sep}$.

3. A Riemann-Roch theorem for metrized complexes of curves

Our aim in this section is to formulate and prove a Riemann-Roch theorem for metrized complexes of κ-curves which generalizes both the classical Riemann-Roch theorem for smooth projective curves over κ and the Riemann-Roch theorem for (metric) graphs from [BN, GK, MZ].

3.1. **Canonical class and statement of Riemann-Roch theorem.** Let $\mathfrak{C}$ be a metrized complex of algebraic curves over κ. Recall that this consists of a metric graph Γ with a model $G = (V, E)$, a weight function $\omega : V \to \mathbb{N}$, and a collection $\{C_v\}_{v \in V}$ of geometrically connected, smooth projective κ-curves of genus $g_v = \omega(v)$ for each vertex $v \in V$. In addition, for each $v \in V$ a collection $\mathcal{A}_v$ of $\deg_G(v)$ points of $C_v(\kappa)$ is given, and these points are labeled with the edges incident to v. The *genus* of $\mathfrak{C}$, denoted $g(\mathfrak{C})$, is by definition $g(\mathfrak{C}) = g(\Gamma) + \sum_{v \in V} g_v$.

Let $\mathcal{W}$ be the divisor $\mathcal{W} = \sum_{v \in V} g_v(v)$ on Γ. Following the notation from [AC], the genus $g(\Gamma, \omega)$ and the canonical class of the (vertex)-weighted metric graph (Γ, ω) are given as follows:

$$g(\Gamma, \omega) = g(\Gamma) + \sum_v g_v = g(\mathfrak{C}), \quad K^{\#} = \sum_v \left(\deg_G(v) + 2g_v - 2\right)(v).$$

For each curve C_v, let K_v denote a divisor of degree $2g_v - 2$ in the canonical class of C_v. Denote by A_v the divisor in C_v consisting of the sum of the $\deg_G(v)$ points in $\mathcal{A}_v$. Let $\mathcal{K}$ be the divisor on $\mathfrak{C}$ defined by

$$\mathcal{K} := K^{\#} \oplus \sum_{v \in V} (K_v + A_v).$$

Note that $\mathcal{K}$ is a well-defined divisor in $\mathrm{Div}(\mathfrak{C})$ since for every $v \in V$, $\deg(K_v + A_v) = 2g_v - 2 + \deg_G(v)$. Note that while the definition of $\mathcal{K}$ depends on the choices of the various canonical divisors K_v in $\mathrm{Div}(C_v)$, the linear equivalence class of $\mathcal{K}$ is independent of these choices. The class of any $\mathcal{K}$ as above is called the *canonical class* of $\mathfrak{C}$.

Remark 3.1. Let X_0 be a strongly semistable curve. The definition of the canonical class $\mathcal{K}$ in $\mathfrak{C}X_0$ given above is compatible with the definition of the dualizing sheaf (also called the sheaf of logarithmic 1-forms) $\omega_{X_0/\kappa}$ [H]. Indeed, if $\alpha : X \to X_0$ denotes the normalization of X_0, and y^e,z^e denote the points of X above the singular point x^e of X_0, for $e \in E(G_0)$, sections the sheaf $\omega_{X_0/\kappa}$ on an open set U consist of all the 1-forms δ on X which are regular everywhere except for possible simple poles at y^e, z^e, with $\mathrm{Res}_{y^e}(\delta) + \mathrm{Res}_{z^e}(\delta) = 0$ for all e. From this, it is easy to see that the restriction of $\omega_{X_0/\kappa}$ to each component X_v of X_0 is the invertible sheaf $\omega_{X_v}(A_v)$ corresponding to $K_v + A_v$ for any canonical divisor K_v on X_v.

For a divisor $\mathcal{D}$ on $\mathfrak{C}$, let $r_{\mathfrak{C}}(\mathcal{D})$ denote the rank of $\mathcal{D}$ as defined in Section 2.

Theorem 3.2 (Riemann-Roch for metrized complexes of algebraic curves). *Let $\mathfrak{C}$ be a metrized complex of κ-curves and $\mathcal{K}$ a divisor in the canonical class of $\mathfrak{C}$. For every divisor $\mathcal{D} \in \mathrm{Div}(\mathfrak{C})$, we have*

$$r_{\mathfrak{C}}(\mathcal{D}) - r_{\mathfrak{C}}(\mathcal{K} - \mathcal{D}) = \deg(\mathcal{D}) - g(\mathfrak{C}) + 1.$$

Let X_0 be a strongly semistable curve of genus g over κ. For any $\mathcal{L} \in \mathrm{Pic}(X_0)$, let $r_c(\mathcal{L})$ be the combinatorial rank of $\mathcal{L}$ as defined in Section 2. As a corollary of Theorem 3.2, Remark 3.1, and Proposition 2.3, we have:

Corollary 3.3 (Riemann-Roch for strongly semistable curves). *For any $\mathcal{L} \in \mathrm{Pic}(X_0)$,*

$$r_c(\mathcal{L}) - r_c(\omega_{X_0/\kappa} \otimes \mathcal{L}^{-1}) = \deg(\mathcal{L}) - g(X_0) + 1.$$

The Riemann-Roch theorem has a number of well-known formal consequences which can be transported to our situation. For example, we obtain the following analogue of Clifford's bound for the rank of a special divisor. (A divisor $\mathcal{D}$ on a metrized complex $\mathfrak{C}$ is called *special* if $K_{\mathfrak{C}} - \mathcal{D}$ is linearly equivalent to an effective divisor.)

Theorem 3.4 (Clifford's theorem for metrized complexes). *For any special divisor $\mathcal{D}$ on a metrized complex $\mathfrak{C}$, $r_{\mathfrak{C}}(\mathcal{D}) \leq \deg(\mathcal{D})/2$.*

Remark 3.5. Our Riemann-Roch theorem seems to be rather different from the classical Riemann-Roch theorem on semistable curves. We refer to [C1] for a discussion of Riemann's and Clifford's theorems for linear series on semistable curves, and for examples showing the failure of Clifford's theorem for $h^0(\mathcal{L})$ with $\mathcal{L} \in \mathrm{Pic}(X_0)$ (even for line bundles in the compactified Picard scheme of X_0).

The rest of this section is devoted to the proof of Theorem 3.2. Our proof follows and extends the original arguments of [BN] in the proof of Riemann-Roch theorem for graphs. In particular, we are going to first extend the notion of reduced divisors to the context of metrized complexes of algebraic curves, and then recall and use an abstract Riemann-Roch criterion from [BN] to derive Theorem 3.2.

3.2. **Reduced divisors.** Let v_0 be a fixed (base) point of Γ. We introduce the notion of v_0-reduced divisors and show that each equivalence class of divisors on $\mathfrak{C}$ contains a quasi-unique v_0-reduced divisor. (We refer to [A, BN, HKN, MZ] for more details on the "classical" notion of reduced divisors on graphs and tropical curves.)

For a closed connected subset S of Γ and a point $v \in \partial S$, the topological boundary of S, the number of "outgoing" tangent directions at v (i.e., tangent directions emanating outward from S) is denoted by $\mathrm{outdeg}_S(v)$; this is also the maximum size of a collection of internally disjoint segments in $\Gamma \setminus (S - \{v\})$ with an end in v. In addition, in the case where v is in V, we denote by $\mathrm{div}_v(\partial S)$ the divisor in C_v associated to the outgoing edges at v; this is by definition the sum of all the points x_v^e of C_v indexed by the edges e leaving S at v. In what follows we refer to a closed connected subset S of Γ as a *cut* in Γ.

Let $\mathcal{D} = D_\Gamma \oplus \sum_{v\in V} D_v$ be a divisor on $\mathfrak{C}$ and S a cut in Γ. A boundary point $x \in \partial S$ is called *saturated* with respect to $\mathcal{D}$ and S if

- $x \notin V$ and $\mathrm{outdeg}_S(x) \le D_\Gamma(x)$; or
- $x = v$ for some $v \in V$ and $D_v - \mathrm{div}_v(\partial S)$ is equivalent to an effective divisor on C_v.

Otherwise, $x \in \partial S$ is called *non-saturated*. A cut S is said to be saturated if all its boundary points are saturated. (When talking about saturated and non-saturated points, we will sometimes omit the divisor $\mathcal{D}$ or the set S if they are clear from the context.)

A divisor $\mathcal{D} = D_\Gamma \oplus \sum_{v\in V} D_v$ on a metrized complex of κ-curves $\mathfrak{C}$ is said to be v_0*-reduced* if the following three properties are satisfied:

- For all points $x \neq v_0$ of Γ, $D_\Gamma(x) \ge 0$, i.e., all the coefficients of D_Γ are non-negative except possibly at the base point v_0.
- For all points $v \in V \setminus \{v_0\}$ ($= V$ if v_0 does not belong to V), there exists an effective divisor E_v linearly equivalent to D_v on C_v.
- For every cut S of Γ which does not contain v_0, there exists a non-saturated point $x \in \partial S$.

We now show that every divisor $\mathcal{D}$ on $\mathfrak{C}$ is linearly equivalent to a quasi-unique v_0-reduced divisor $\mathcal{D}^{v_0} = D_\Gamma^{v_0} \oplus \sum_{v\in V} D_v^{v_0}$. The quasi-uniqueness is understood in the following sense: $D_\Gamma^{v_0}$ is unique on Γ and for all $v \in V$, the divisor class $[D_v^{v_0}]$ on C_v defined by $D_v^{v_0}$ is unique.

Theorem 3.6. *Let $\mathfrak{C}$ be a metrized complex of κ-curves and v_0 a base point of Γ. For every divisor $\mathcal{D} \in \mathrm{Div}(\mathfrak{C})$, there exists a quasi-unique v_0-reduced divisor $\mathcal{D}^{v_0}$ such that $\mathcal{D}^{v_0} \sim \mathcal{D}$.*

Proof. We follow the lines of the proof given in [A] for the existence of v_0-reduced divisors on tropical curves, and adapt it to the context of metrized complexes of κ-curves.

Let $\mathcal{T}_\mathcal{D}$ be the set of all divisors $\mathcal{D}' = D'_\Gamma \oplus \sum_v D'_v$ linearly equivalent to $\mathcal{D}$ such that:

- All the coefficients of D'_Γ are non-negative at every point of Γ except possibly at v_0.
- For each point $v \in V \setminus \{v_0\}$, D_v has non-negative rank on C_v.
- The coefficient of v_0 in D'_Γ is maximal with respect to the two properties above.

Note that $\mathcal{T}_\mathcal{D}$ is non-empty. Indeed, the arguments from [A] can be easily adapted to show that there is a divisor $\mathcal{D}' = D'_\Gamma \oplus \sum_{v\in V} D'_v$ linearly equivalent to $\mathcal{D}$ on $\mathfrak{C}$ with the property that the coefficient of every point $x \in \Gamma \setminus (V \cup \{v_0\})$ is non-negative and the coefficient of D'_Γ at each point $v \in V$ (which coincides with the degree of D'_v) is at least g_v.

By the (classical) Riemann-Roch theorem, one infers that each D'_v has non-negative rank on C_v.

Define $T_{\mathcal{D}}$ as the set of all $D'_\Gamma \in \mathrm{Div}(\Gamma)$ for which there exists $\mathcal{D}' \in \mathcal{T}_{\mathcal{D}}$ with $\mathcal{D}' = D'_\Gamma \oplus \sum_{v\in V} D'_v$. Each divisor $D'_\Gamma \in T_{\mathcal{D}}$ has degree $\deg(\mathcal{D})$, and all divisors in $T_{\mathcal{D}}$ have the same coefficient at v_0. Thus the number

$$N := \sum_{x\in\Gamma\setminus\{v_0\}} D'_\Gamma(x)$$

is independent of the choice of a divisor D'_Γ in $T_{\mathcal{D}}$. As a consequence, $T_{\mathcal{D}}$ inherits a natural topology from the topology of Γ, since it can be embedded as a subset of $\mathrm{Sym}^N\Gamma$. Since Γ is compact, one sees easily that $T_{\mathcal{D}}$ is compact as well.

Let $\mathcal{A}$ be the subset of $\mathbb{R}^N$ defined by

$$\mathcal{A} := \Big\{(x_1,\dots,x_N)\,|\,0\le x_1\le x_2\le\cdots\le x_N\Big\}.$$

Note that $\mathcal{A}$ has a natural total order defined by the lexicographical rule $(x_i)\le(y_i)$ iff $x_1=y_1,\dots,x_i=y_i$ and $x_{i+1}<y_{i+1}$.

Define a continuous map $F: T_{\mathcal{D}} \to \mathcal{A}$ as follows. For each divisor $D'_\Gamma \in T_{\mathcal{D}}$, consider the multiset $A(D'_\Gamma)$ of points in $\Gamma\setminus\{v_0\}$ where each point $v\ne v_0$ appears in this multiset exactly $D'_\Gamma(v)$ times. Define $F(D'_\Gamma)$ to be the point of $\mathcal{A}$ defined by the multiset of distances $\mathrm{dist}_\Gamma(v,v_0)$ for $v\in A(D'_\Gamma)$, ordered in an increasing way. One verifies easily, as in [A], that the map F is continuous. Since $T_{\mathcal{D}}$ is compact and F is continuous, there exists a divisor $D^{v_0}_\Gamma$ in $T_{\mathcal{D}}$ such that F takes its minimum value at $D^{v_0}_\Gamma$, i.e., $F(D^{v_0}_\Gamma)=\min_{D'_\Gamma\in T_{\mathcal{D}}} F(D'_\Gamma)$. Let $\mathcal{D}^{v_0}$ be a point of $\mathcal{T}_{\mathcal{D}}$ giving rise to $D^{v_0}_\Gamma$ in $T_{\mathcal{D}}$, i.e., such that $\mathcal{D}^{v_0} = D^{v_0}_\Gamma \oplus \sum_{v\in V} D^{v_0}_v$.

We claim that $\mathcal{D}^{v_0}\sim\mathcal{D}$ is v_0-reduced. The first two properties are clearly satisfied: all the coefficients $D^{v_0}_\Gamma(x)$ are non-negative for $x\ne v_0$ in Γ, and D_v is linearly equivalent to an effective divisor on C_v for all $v\in V$. The only thing we need to check is that every cut S which does not contain v_0 has a non-saturated point on its boundary. For the sake of contradiction, suppose this is not the case, and let S be a closed connected set violating this condition. This means that the following hold:

- For all $x\in\partial S\setminus V$, $\mathrm{outdeg}_S(x)\le D^{v_0}_\Gamma(v)$.
- For each $v\in\partial S\cap V$, $D^{v_0}_v - \mathrm{div}_v(\partial S)$ has non-negative rank on C_v, i.e., there exists a rational function f_v on C_v such that $D^{v_0}_v - \mathrm{div}_v(\partial S)+\mathrm{div}(f_v)\ge 0$.

By the definition of outdeg_S, there exists an $\epsilon>0$ such that for each vertex $x\in\partial S$, there are closed segments $I^x_1,\dots,I^x_{\mathrm{outdeg}_S(x)}$ emanating from x with the following properties:

- For $x\in\partial S$ and $1\le j\le\mathrm{outdeg}_S(x)$, the half-open segments $I^x_j\setminus x$ are disjoint from S and from each other and do not contain v_0.
- Each segment I^x_j has length ϵ, for $x\in\partial S$ and $1\le j\le\mathrm{outdeg}_S(x)$.

Define a rational function $f_\Gamma:\Gamma\to\mathbb{R}$ as follows: f_Γ is identically zero on S, is linear of slope -1 on each interval I^x_j for $x\in\partial S$ and $1\le j\le\mathrm{outdeg}_S(x)$, and takes the constant value $-\epsilon$ at all points of $\Gamma\setminus\Big(S\cup\bigcup_{x,j}I^x_j\Big)$. Let $\mathfrak{f}$ be the rational function on $\mathfrak{C}$ consisting of the rational function f_Γ on Γ and the rational functions f_v on C_v. (Recall that f_v is chosen so that $D^{v_0}_v-\mathrm{div}_v(\partial S)+\mathrm{div}(f_v)\ge 0$.)

Consider the divisor $\mathcal{D}^* = \mathcal{D}^{v_0}+\mathrm{div}(\mathfrak{f})$ on $\mathfrak{C}$, and write $\mathcal{D}^* = D^*_\Gamma\oplus\sum_{v\in V}D^*_v$. We claim that $\mathcal{D}^*$ and D^*_Γ lie in $\mathcal{T}_{\mathcal{D}}$ and $T_{\mathcal{D}}$, respectively. First of all, since for each point

$x \in \Gamma \setminus (V \cup \{v_0\})$ we have $\mathrm{outdeg}_S(x) \leq D_\Gamma^{v_0}(x)$, it is clear that $D_\Gamma^{v_0} + \mathrm{div}(f_\Gamma)$ has non-negative coefficient at x. Secondly, each D_v^* is an effective divisor on C_v and thus $\deg(D_v^*) = D_\Gamma^*(v) \geq 0$ for each $v \in V$. In addition, $D_\Gamma^*(v_0) = D_\Gamma^{v_0}(v_0)$. This proves that $\mathcal{D}^* \in \mathcal{T}_\mathcal{D}$, and thus that $D_\Gamma^* \in T_\mathcal{D}$. Since the distance from S to v_0 is strictly larger than the distance from the set $S \cup \bigcup_{x\in\partial S} \bigcup_{j\leq \mathrm{outdeg}_S(x)} I_j^x$ to v_0, and since the boundary points of $S \cup \bigcup_{x\in\partial S} \bigcup_{j\leq \mathrm{outdeg}_S(x)} I_j^x$ are all in the multiset $A(D_\Gamma^*)$, it follows easily that $F(D_\Gamma^*) < F(D_\Gamma^{v_0})$, contradicting the choice of $D_\Gamma^{v_0}$. Thus the divisor $\mathcal{D}^{v_0}$ is v_0-reduced, which proves the existence part of the theorem.

It remains to prove the quasi-uniqueness. Assume for the sake of contradiction that there are linearly equivalent v_0-reduced divisors $\mathcal{D} = D_\Gamma \oplus \sum_{v\in V} D_v$ and $\mathcal{D}' = D'_\Gamma \oplus \sum_{v\in V} D'_v$ on $\mathfrak{C}$ such that either $D_\Gamma \neq D'_\Gamma$ or D_v and D'_v are not linearly equivalent on C_v for some $v \in V$. Then there exists a non-constant rational function $\mathfrak{f} = \{f_\Gamma; \{f_v\}\}$ on $\mathfrak{C}$ such that $\mathcal{D}' = \mathcal{D} + \mathrm{div}(\mathfrak{f})$. If f_Γ is constant then we obtain an immediate contradiction, so we may assume that f_Γ is non-constant. Without loss of generality, we may assume that f_Γ does not take its maximum at v_0 (otherwise, we can interchange the role of $\mathcal{D}$ and $\mathcal{D}'$). Let S be a connected component of the set of all points where f_Γ takes its maximum. Note that $v_0 \notin S$. For all point $x \in \partial S$, the slope of f_Γ at any outgoing segment I_j^x emanating from v is at most -1. Since $\mathcal{D}$ is v_0-reduced, there exists a point $x \in \partial S$ such that either $x \notin V$ and $D_\Gamma(x) < \mathrm{outdeg}_S(x)$, or $x = v$ for a point $v \in V$ and the divisor $D_v - \mathrm{div}(\partial S)$ has negative rank. In the first case, $D'_\Gamma(x) \leq D_\Gamma(x) - \mathrm{outdeg}_S(x) < 0$, contradicting the assumption that the coefficient of D'_Γ is non-negative at $x \neq v_0$. In the second case, $D'_v \sim D_v + \sum_{e\in E:\, e\sim v} \mathrm{slp}_e(f_\Gamma)(x_v^e) \leq D_v - \mathrm{div}_v(\partial S)$, which implies that D'_v has negative rank, a contradiction. □

3.3. **Riemann-Roch criterion.** We recall the abstract Riemann-Roch criterion from [BN]. Let X be a non-empty set and let $\mathrm{Div}(X)$ be the free abelian group on X. Elements of $\mathrm{Div}(X)$ are called *divisors* on X, and divisors E with $E \geq 0$ are called *effective.* The degree of a divisor D is the sum of the coefficients in D of elements of X. Let $\sim$ be an equivalence relation on $\mathrm{Div}(X)$ satisfying the following two properties:

(E1) If $D \sim D'$, then $\deg(D) = \deg(D')$.
(E2) If $D_1 \sim D'_1$ and $D_2 \sim D'_2$, then $D_1 + D'_1 \sim D_2 + D'_2$.

For $D \in \mathrm{Div}(X)$ we define $|D|$ to be the set of all effective divisors $E \sim D$, and we define $r(D)$ to be

$$r(D) = \min_{E\geq 0:|D-E|=\emptyset} \deg(E) - 1.$$

Let g be a non-negative integer, and define

$$\mathcal{N} = \{D \in \mathrm{Div}(X) \ : \ \deg(D) = g - 1 \text{ and } |D| = \emptyset\}.$$

Finally, let K be an element of $\mathrm{Div}(X)$ having degree $2g - 2$. The following theorem from [BN] gives necessary and sufficient conditions for the Riemann-Roch formula to hold for elements of $\mathrm{Div}(X)/\sim$ and for the rank function r defined above. Define $\epsilon : \mathrm{Div}(X) \to \mathbb{Z}/2\mathbb{Z}$ by setting $\epsilon(D) = 0$ if $|D| \neq \emptyset$ and $\epsilon(D) = 1$ if $|D| = \emptyset$.

Theorem 3.7. *The Riemann-Roch formula*

$$r(D) - r(K - D) = deg(D) + 1 - g$$

holds for all $D \in \mathrm{Div}(X)$ if and only if the following two properties are satisfied:

(RR1) *For every $D \in \mathrm{Div}(X)$, there exists $\mathfrak{n} \in \mathcal{N}$ such that $\epsilon(D) + \epsilon(\mathfrak{n} - D) = 1$.*

(RR2) *For every* $D \in \mathrm{Div}(X)$ *with* $\deg(D) = g - 1$, $\epsilon(D) + \epsilon(K - D) = 0$.

The divisors in $\mathcal{N}$ are called *maximal non-special.*

Remark 3.8. Recall that a divisor D on a smooth projective curve C over κ is called *special* if $r(K_C - D) \geq 0$. By Riemann-Roch, a divisor of degree at least $g(C)$ on C can never be special, and if $\deg(D) = g(C) - 1$ then D is special if and only if $r(D) \geq 0$. This explains the term "maximal non-special divisor'.

3.4. **Description of maximal non-special divisors.** We now provide an explicit description of the maximal non-special divisors on metrized complexes of algebraic curves, generalizing the corresponding description of maximal non-special divisors on graphs and metric graphs from [BN, MZ, HKN]. Let $\mathfrak{C}$ be a metrized complex of algebraic curves. For all $v \in V$, let $\mathcal{N}_v$ be the set of all maximal non-special divisors on C_v.

An *acyclic orientation* on Γ is an acyclic orientation on some model of Γ, i.e., a decomposition of Γ into closed directed edges with disjoint interiors such that no directed cycles are formed. Given an acyclic orientation π on Γ, we denote by $\deg^+_\pi(x)$ the number of tangent directions emanating from x which are compatible with π. Note that for all but finitely many points of Γ, $\deg^+_\pi(x) = 1$. For $v \in V$, we denote by $E^+(v)$ the set of edges incident to v which (locally near v) are oriented outward from v.

Given a collection of maximal non-special divisors $D_v \in \mathcal{N}_v$ for each $v \in V$, together with an acyclic orientation π of Γ, define a corresponding divisor $\mathcal{D}^{\pi,\{D_v\}} = D^\pi_\Gamma \oplus \Sigma_{v\in V} D^\pi_v \in \mathrm{Div}(\mathfrak{C})$ as follows. First, set

$$D^\pi_\Gamma := \sum_{x\in\Gamma} (\deg^+_\pi(x) + g_x - 1)(x).$$

(Recall that g_x is zero if $x \notin V$ and is the genus of C_x for $x \in V$.)

Next, for each point v of V, let $\mathcal{A}^\pi_v$ be the set of all points in $\mathcal{A}_v$ indexed by an oriented edge of π with v as the tail. Define the divisor A^π_v on C_v to be the sum of the elements of $\mathcal{A}^\pi_v$, and define $D^\pi_v := A^\pi_v + D_v$. Note that $\deg(D^\pi_v) = \deg^+_\pi(v) + g_v - 1 = D^\pi_\Gamma(v)$. The degree of $\mathcal{D}^{\pi,\{D_v\}}$ is given by

$$\begin{aligned}\deg(\mathcal{D}^{\pi,\{D_v\}}) &= \sum_{v\in\widetilde{V}} \big(\deg^+_\pi(v) + g_v - 1\big)\\ &= \Big[\sum_{v\in\widetilde{V}} \big(\deg^+_\pi(v) - 1\big)\Big] + \sum_{v\in V} g_v = g(\Gamma) - 1 + \sum_v g_v = g(\mathfrak{C}) - 1.\end{aligned}$$

A divisor $\mathcal{M} \in \mathrm{Div}(\mathfrak{C})$ of the form $\mathcal{D}^{\pi,\{D_v\}} = D^\pi_\Gamma \oplus \sum D^\pi_v$ is called a *moderator* on $\mathfrak{C}$. (This terminology comes from [MZ].) Given a moderator $\mathcal{M} = \mathcal{D}^{\pi,\{D_v\}}$, the *dual moderator* $\bar{\mathcal{M}}$ is defined to be $\mathcal{D}^{\bar\pi,\{K_v - D_v\}}$, where $\bar\pi$ is obtained from π by reversing the orientation of every oriented segment. It is easy to see that $\mathcal{M} + \bar{\mathcal{M}}$ belongs to the canonical class on $\mathfrak{C}$.

The following two lemmas are essentially what we need for the proof of Theorem 3.2.

Lemma 3.9. *For any acyclic orientation* π *of* Γ *and any collection* $D_v \in \mathcal{N}_v$ *of maximal non-special divisors on* C_v*, the moderator* $\mathcal{D}^{\pi,\{D_v\}}$ *is a maximal non-special divisor on* $\mathfrak{C}$.

Proof. Suppose for the sake of contradiction that there is a rational function $\mathfrak{f} = \{f_\Gamma; \{f_v\}\}$ on $\mathfrak{C}$ such that $\mathcal{D}^{\pi,\{D_v\}} + \mathrm{div}(\mathfrak{f}) \geq 0$. Let $\widetilde{G} = (\widetilde{V}, \widetilde{E})$ be a (loopless) model of Γ such that the orientation π of Γ is induced by an acyclic orientation $\widetilde{G}^\pi$ on the edges of $\widetilde{G}$. Let v_0 be a sink vertex of the directed graph $\widetilde{G}^\pi$ (i.e., a vertex for which the out-degree is zero). Let S be the closed subset of Γ consisting of all the points of Γ where f_Γ takes its maximum value. We claim that $v_0 \notin S$. Indeed, if f_Γ achieves its maximum value at v_0, then:

- If $v_0 \notin V$, the coefficient of v_0 in D^π_Γ is at most -1, a contradiction.
- If $v_0 \in V$, then $D_{v_0} + \mathrm{div}(f_{v_0}) \geq D^\pi_{v_0} + \mathrm{div}_{v_0}(f_\Gamma) + \mathrm{div}(f_{v_0}) \geq 0$, contradicting the assumption that D_{v_0} is a maximal non-special divisor on C_{v_0}.

We now claim that $\partial S \subseteq \widetilde{V}$. Indeed, if there was a point $x \in \partial S \setminus \widetilde{V}$, then since the support of D^π_Γ is contained in $\widetilde{V}$ we would obtain $D^\pi_\Gamma(x) + \mathrm{div}_x(f_\Gamma) < 0$, contradicting the assumptions on $\mathfrak{f}$.

Next, we claim that there exists a point $u \in \partial S$ such that $\deg^+_\pi(u) = \mathrm{outdeg}_S(u)$. To see this, consider the restriction of π to the subgraph $\widetilde{G}[\widetilde{V} \cap S]$ of $\widetilde{G}$ induced by $\widetilde{V} \cap S \subset \widetilde{V}$. The resulting directed graph $\widetilde{G}^\pi[\widetilde{V} \cap S]$ is acyclic and therefore contains a sink vertex u. But for any vertex $v \in \widetilde{V} \cap S \setminus \partial S$, the out-degrees of v in $\widetilde{G}^\pi$ and $\widetilde{G}^\pi[\widetilde{V} \cap S]$ are the same. It follows easily that $\deg^+_\pi(u) = \mathrm{outdeg}_S(u)$, which proves our claim.

By the choice of $u \in \partial S$ and by the definition of S, f_Γ has strictly negative slope along all out-going edges at u, and it has has slope zero along all other edges incident to u. This shows that $\mathrm{ord}_u(f_\Gamma) \leq -\deg^+_\pi(u)$. There are now two cases:

- $u \in \widetilde{V} \setminus V$, in which case $D^\pi_\Gamma(u) + \mathrm{ord}_u(f_\Gamma) \leq \deg^+_\pi(u) - 1 - \deg^+_\pi(u) < 0$, contradicting the choice of the rational function $\mathfrak{f}$ on $\mathfrak{C}$.
- $u \in V$, in which case
$$\mathcal{D}^{\pi,\{D_v\}} + \mathrm{div}(\mathfrak{f}) = (D^\pi_\Gamma + \mathrm{div}(f_\Gamma)) \oplus \sum_v \Big[D^\pi_v + \mathrm{div}(f_v) + \sum_{e \sim v} \mathrm{slp}_e(f_\Gamma)(x^e_v) \Big].$$
Once again, for each edge e incident to u, we have $\mathrm{slp}_e(f_\Gamma) \leq 0$, with strict inequality if e is an out-going edge at u in $\widetilde{G}^\pi$. By the definition of D^π_u, we infer that $D^\pi_u + \sum_{e \in \widetilde{E}: e \sim u} \mathrm{slp}_e(f_\Gamma)(x^e_u) \leq D_u$, which shows that $D^\pi_u + \sum_{e \in \widetilde{E}: e \sim u} \mathrm{slp}_e(f_\Gamma)(x^e_u) + \mathrm{div}(f_u)$ cannot be effective. This contradicts the hypothesis we made on $\mathfrak{f}$.

□

For $v_0 \in \Gamma$, we denote by $\mathcal{AO}_{v_0}(\Gamma)$ the set of all acyclic orientations of Γ with a unique sink at v_0, i.e., such that v_0 has out-degree zero and all other points of Γ have out-degree at least one.

Lemma 3.10. *Let $\mathcal{D} = D_\Gamma \oplus \sum_{v \in V} D_v$ be a v_0-reduced divisor on $\mathfrak{C}$. Then $\epsilon(\mathcal{D}) = 0$ if and only if the following two conditions hold:*

- *The coefficient $D_\Gamma(v_0)$ of D_Γ at v_0 is non-negative.*
- *If $v_0 \in V$, the divisor D_{v_0} on C_{v_0} has non-negative rank.*

More precisely, if the above conditions do not both hold, then there exists a moderator $\mathcal{M} = \mathcal{D}^{\pi,\{D^*_v\}}$ with $\pi \in \mathcal{AO}_{v_0}(\Gamma)$ such that $\mathcal{D} \leq \mathcal{M}$.

Proof. First suppose that $D_\Gamma(v_0) \geq 0$, and in the case $v_0 \in V$ that there exists a rational function f_{v_0} on C_{v_0} such that $D_{v_0} + \mathrm{div}(f_{v_0}) \geq 0$. We infer that $D_\Gamma \geq 0$, and since each D_v for $v \in V \setminus \{v_0\}$ is of non-negative rank, there exists a rational function f_v on C_v such that $D_v + \mathrm{div}(f_v) \geq 0$. Let $\mathfrak{f}$ be a rational function on $\mathfrak{C}$ consisting of a constant rational function on Γ and f_v on C_v for $v \in V$. We obviously have $\mathcal{D} + \mathrm{div}(\mathfrak{f}) \geq 0$, which implies that $\epsilon(\mathcal{D}) = 0$.

For the other direction, assume that $D_\Gamma(v_0) < 0$ if $v_0 \notin V$, and that D_{v_0} has negative rank if $v_0 \in V$. To show that $\epsilon(\mathcal{D}) = 1$, by Lemma 3.9 it will be enough to show the existence of an acyclic orientation $\pi \in \mathcal{AO}_{v_0}(\Gamma)$ and a set of maximal non-special divisors $D_v^* \in \mathcal{N}_v$ such that $\mathcal{D} \leq \mathcal{D}^{\pi,\{D_v^*\}}$.

Let $\widetilde{V}$ be the union of V and all the points in the support of D_Γ. Note that in both case above, $v_0 \in \widetilde{V}$. Let $\widetilde{G} = (\widetilde{V}, \widetilde{E})$ be the corresponding model of Γ. We are going to recursively define an orientation π of $\widetilde{G}$ and the collection $\{D_v^*\}$ by handling at each step the orientation of all the edges incident to a vertex $v \in \widetilde{V}$ and the maximal non-special divisor D_v^* in the case $v \in V$. We start by considering the vertex $v_0 \in \widetilde{V}$. The orientations of all the edges incident to v_0 are defined so that $\deg_\pi^+(v_0) = 0$; in other words, all these oriented edges are incoming at v_0. Note that in the case $v_0 \notin V$, we have $D_\Gamma(v_0) \leq -1 = D_\Gamma^\pi(v_0)$. In the case $v_0 \in V$, since $r(D_{v_0}) < 0$ it follows from the Riemann-Roch theorem for C_{v_0} that there exists a maximal non-special divisor $D_{v_0}^*$ such that $D_{v_0} \leq D_{v_0}^*$.

Suppose that the orientation of all edges adjacent to vertices $v_0, \dots, v_i \in \widetilde{V}$ has been defined, and that for all $v_j \in V$ with $j \leq i$, a maximal non-special divisor D_v^* on C_v has been given. Let S_i be a connected component of the induced subgraph $\widetilde{G}[\widetilde{V} \setminus \{v_0, \dots, v_i\}]$. Note that S_i defines a cut in Γ (which by abuse of the notation we continue to denote by S_i) not containing v_0. Since $\mathcal{D}$ is v_0-reduced, there exists a point v_{i+1} on the boundary of S_i which is non-saturated. (Note that v_{i+1} also lies in $\widetilde{V}$.) This means that either:

1. $v_{i+1} \in \Gamma \setminus V$ and $\mathrm{outdeg}_{S_i}(v_{i+1}) > D_\Gamma(v_{i+1})$; or
2. $v_{i+1} \in V$ and $D_v - \mathrm{div}_v(\partial S_i)$ has negative rank.

All out-going edges from S_i adjacent to v_{i+1} have already been oriented (and are out-going from v_{i+1} by the definition of the orientation π). Orient all the other edges incident to v_{i+1} in such a way that they are all in-coming at v_{i+1}. Note that $\deg_\pi^+(v_{i+1}) = \mathrm{outdeg}_{S_i}(v_{i+1})$, and so in Case (1),

$$D_\Gamma(v_{i+1}) \leq \deg_\pi^+(v_{i+1}) - 1.$$

In Case (2), since $r(D_v - \mathrm{div}_v(\partial S_i)) = -1$, it follows from the Riemann-Roch theorem for C_v that there exists a maximal non-special divisor D_v^* such that

$$D_v - \mathrm{div}_v(\partial S_i) \leq D_v^*.$$

This shows that $D_v \leq \mathrm{div}_v(\partial S_i) + D_v^*$. By the definition of the orientation π we must have $D_{v+1}^{*\,\pi} = \mathrm{div}_{v_{i+1}}(\partial S_i) + D_v^*$. Therefore, by the definition of the orientation π at v_{i+1} and the choice of $D_{v_{i+1}}^*$, we have $D_{v_{i+1}} \leq D_{v_{i+1}}^{*\,\pi}$.

Let π be the orientation of Γ just constructed, and let $\{D_v^*\}_{v \in V}$ be the collection maximal non-special divisors on C_v defined above. By the definition of π, the only vertex of $\widetilde{V}$ with out-degree zero is v_0 and thus π belongs to $\mathcal{AO}_{v_0}(\Gamma)$. We clearly have $\mathcal{D} \leq \mathcal{D}^{\pi,\{D_v^*\}}$ as well, which completes the proof. □

As a corollary of the above lemma, we obtain:

Corollary 3.11. *Any maximal non-special divisor $\mathcal{D}$ on $\mathfrak{C}$ is linearly equivalent to a moderator.*

Proof. Fix $v_0 \in \Gamma$. Let $\mathcal{D}$ be a maximal non-special divisor on $\mathfrak{C}$ and let $\mathcal{D}^{v_0} = D_\Gamma^{v_0} \oplus \sum_{v \in V} D_v^{v_0}$ be the v_0-reduced divisor linearly equivalent to $\mathcal{D}$. By Lemma 3.10, there exists a moderator $\mathcal{M}$ such that $\mathcal{D}^{v_0} \leq \mathcal{M}$. The effective divisor $\mathcal{D}^{\pi,\{D_v^*\}} - \mathcal{D}^{v_0}$ has degree zero, and so must be equal to zero. $\square$

3.5. **Proof of Theorem 3.2.** We will show that Conditions (RR1) and (RR2) in the Riemann-Roch criterion (Theorem 3.7) are verified if we choose $\mathcal{K}$ in the canonical class of $\mathfrak{C}$ and let g be the genus of $\mathfrak{C}$. In our application of Theorem 3.7, we take X to be the union of $\Gamma \setminus V$ and $C_v(\kappa)$ for $v \in V$, c.f. Remark 2.1.

Proof of (RR1). Let $\mathcal{D} = D_\Gamma \oplus \sum_{v\in V} D_v$ be a divisor on $\mathfrak{C}$, which we can assume to be v_0-reduced for some $v_0 \notin V$. The proof breaks into two cases, depending on whether $D_\Gamma(v_0) \geq 0$ or $D_\Gamma(v_0) < 0$:

- If $D_\Gamma(v_0) \geq 0$, then by Lemma 3.10 we have $\epsilon(\mathcal{D}) = 0$ and there exists an effective divisor $\mathcal{E}$ on $\mathfrak{C}$ such that $\mathcal{D} \sim \mathcal{E}$. For each maximal non-special divisor $\mathfrak{n} \in \mathcal{N}$, we have $\epsilon(\mathfrak{n} - \mathcal{D}) = \epsilon(\mathfrak{n} - \mathcal{E}) = 1$. This shows that $\epsilon(\mathcal{D}) + \epsilon(\mathfrak{n} - \mathcal{D}) = 1$.
- If $D_\Gamma(v_0) < 0$, then by Lemma 3.10 we have $\epsilon(\mathcal{D}) = 1$, and in addition there exists a moderator $\mathcal{M}$ such that $\mathcal{D} \leq \mathcal{M}$, i.e., $\epsilon(\mathcal{M} - \mathcal{D}) = 0$. By Lemma 3.9, we deduce the existence of a maximal non-special divisor $\mathcal{M} \in \mathcal{N}$ such that $\epsilon(\mathcal{M} - \mathcal{D}) + \epsilon(\mathcal{D}) = 1$.

Proof of (RR2). Let $\mathcal{D} = D_\Gamma \oplus \sum_{v\in V} D_v$ be a divisor of degree $g(\mathfrak{C}) - 1$ on $\mathfrak{C}$. By symmetry, to prove (RR2) we need only show that if $\epsilon(\mathcal{D}) = 1$ then $\epsilon(\mathcal{K} - \mathcal{D}) = 1$. If $\epsilon(\mathcal{D}) = 1$ then $\mathcal{D}$ is maximal non-special, so by Corollary 3.11 there exists a moderator $\mathcal{M}$ on $\mathfrak{C}$ such that $\mathcal{D} \sim \mathcal{M}$. The dual moderator $\bar{\mathcal{M}}$ has the property that $\bar{\mathcal{M}} + \mathcal{M}$ is in the canonical class of $\mathfrak{C}$. We obtain $\epsilon(\mathcal{K} - \mathcal{M}) = 1$, and thus $\epsilon(\mathcal{K} - \mathcal{D}) = 1$ as desired.

The proof of Theorem 3.2 is now complete.

4. Background on Berkovich curves

In this section we provide a brief overview of the parts of Berkovich's theory of non-Archimedean analytic spaces which are important for our purposes. We concentrate mainly on the case of curves, using [BPR, Section 5] as our primary reference.

4.1. **Non-Archimedean analytic spaces.** Let $\mathbb{K}$ be a complete and algebraically closed non-trivially valued non-Archimedean field with valuation ring R and residue field κ. Let $X/\mathbb{K}$ be an algebraic variety, i.e., a reduced, separated scheme of finite type over $\mathbb{K}$. The (Berkovich) *analytification* X^{an} of X, as a topological space, can be defined as follows. Points of X^{an} can be identified with pairs $x = (\xi, |\,|_\xi)$ consisting of a (scheme-theoretic) point $\xi \in X$ and an extension $|\,|_\xi$ of the absolute value on $\mathbb{K}$ to the residue field $\mathbb{K}(\xi)$ of ξ. The topology on X^{an} is the weakest one for which $U^{\mathrm{an}} \subset X^{\mathrm{an}}$ is open for every open affine subset $U \subset X$ and the function $x \mapsto |f(x)|$ is continuous for every $f \in \mathcal{O}_X(U)$. (Here, by a standard abuse of notation, $f(x)$ denotes the image of f in the residue field of ξ, where $x = (\xi, |\,|_\xi)$, and $|f(x)|$ is shorthand for $|f(x)|_\xi$.) There is a canonical inclusion from the set $X(\mathbb{K})$ of closed points of X into X^{an} with dense image. The induced topology on $X(\mathbb{K})$ is the canonical analytic topology with respect to which $X(\mathbb{K})$ is totally disconnected and not locally compact. By contrast, the larger space X^{an} is locally compact, Hausdorff, and locally path-connected. Furthermore, X^{an} is compact iff X is proper and path-connected iff X is connected. There is a category of (Berkovich) $\mathbb{K}$-analytic spaces which includes as a special case the analytifications of algebraic varieties over $\mathbb{K}$. The association $X \mapsto X^{\mathrm{an}}$ is a functor from algebraic varieties over $\mathbb{K}$ to $\mathbb{K}$-analytic spaces. Any open subset X^{an} can be given a $\mathbb{K}$-analytic structure in a natural way.

4.2. **Open discs and open annuli.** Points x in the Berkovich affine line $(\mathbf{A}^1)^{\mathrm{an}}$ over $\mathbb{K}$ can be identified with multiplicative seminorms $|\,|_x$ on the one-variable polynomial ring $\mathbb{K}[T]$ extending the given absolute value on $\mathbb{K}$. An *open ball* in $(\mathbf{A}^1)^{\mathrm{an}}$ is an open subset of the form $\{x \in (\mathbf{A}^1)^{\mathrm{an}} \mid |T|_x < R\}$ for some $R > 0$ (with the induced analytic structure). Similarly, a *punctured open ball* is an open subset of the form $\{x \in (\mathbf{A}^1)^{\mathrm{an}} \mid 0 < |T|_x < R\}$ and an *open annulus* is an open subset of the form $\{x \in (\mathbf{A}^1)^{\mathrm{an}} \mid r < |T|_x < R\}$ for some $0 < r < R$. A *generalized open annulus* is either a punctured open ball (considered as an open annulus with $r = 0$) or an open annulus. Each generalized open annulus $\mathbf{A}$ contains a canonical *skeleton* $\Sigma = \Sigma(\mathbf{A})$, which is the set of all multiplicative norms $|\,|_x$ on $\mathbb{K}[T]$ of the form $|\sum a_i T^i|_x = \max\{|a_i|\rho^i\}$ with $r < \rho < R$. As topological spaces, Σ is homeomorphic to an open interval and there is a canonical retraction map from $\mathbf{A}$ onto Σ. There is also a *canonical metric* on Σ: if $\mathbf{A}$ is an open annulus, we identify Σ with the open interval $(\log r, \log R)$, and if $\mathbf{A}$ is an punctured open ball then we identify Σ with the open interval $(-\infty, \log R)$.

4.3. **Types of points.** For the rest of §4, let $X/\mathbb{K}$ be a smooth, proper, connected algebraic curve. Points of X^{an} are traditionally classified into four types, according to algebraic properties of their completed residue fields. The *completed residue field* $\mathcal{H}(x)$ of $x = (\xi, |\,|_\xi) \in X^{\mathrm{an}}$ is the completion of $\mathbb{K}(\xi)$ with respect to $|\,|_\xi$. We define $s(x)$ to be the transcendence degree of the residue field of $\mathcal{H}(x)$ over the residue field of $\mathbb{K}$, and $t(x)$ to be the rank of the finitely generated abelian group $|\mathcal{H}(x)^\times|/|\mathbb{K}^\times|$. By *Abhyankar's inequality*, $s(x) + t(x) \leq 1$. We say that x is of *type 1* if $\mathcal{H}(x) \cong \mathbb{K}$, of *type 2* if $s(x) = 1$, of *type 3* if $t(x) = 1$, and of *type 4* if $s(x) = t(x) = 0$ but $\mathcal{H}(x) \ncong \mathbb{K}$. The points of type 1 in X^{an} are precisely the points of $X(\mathbb{K})$.

4.4. **Semistable vertex sets and skeleta.** A *semistable vertex set* for X^{an} is a finite set V of type 2 points of X^{an} such that the complement of V in X^{an} is isomorphic (as a $\mathbb{K}$-analytic space) to the disjoint union of a finite number of open annuli and an infinite number of open balls. (Such a disjoint union is called the *semistable decomposition* of X^{an} associated to V.) It follows from the *semistable reduction theorem* that semistable vertex sets always exist, and more generally that any finite set of type 2 points of X^{an} is contained in a semistable vertex set. The *skeleton* $\Gamma = \Sigma(X^{\mathrm{an}}, V)$ of X^{an} with respect to a semistable vertex set V is the union (inside X^{an}) of V and the skeletons of each of the open annuli in the semistable decomposition $\Delta(V)$ associated to V. Using the canonical metric on the skeletons of these open annuli, Γ can be naturally viewed as a (finite) *metric graph* contained in X^{an}.

One can define in a similar way semistable vertex sets and skeleta for an affine curve X'. In this case, one must also allow a finite number of punctured open balls in the semistable decomposition and the skeleton is a topologically finite but not necessarily finite length metric graph – it will contain a finite number of infinite rays corresponding to the points of $X \setminus X'$, where X is the projective completion of X'.

The skeleton Γ of a semistable vertex set V comes equipped with a natural model G whose vertices are the points of V and whose edges correspond bijectively to the open annuli in the semistable decomposition associated to V. A semistable vertex set V is called *strongly semistable* if the graph G has no loop edges. Every semistable vertex set is contained in a strongly semistable vertex set.

There is a canonical retraction map $\tau = \tau_V : X^{\mathrm{an}} \twoheadrightarrow \Gamma$ which sends a point $x \in X^{\mathrm{an}}$ to itself (if $x \in V$), to the retraction of x to $\Sigma(\mathbf{A})$ (if x belongs to an open annulus $\mathbf{A}$ in

$\Delta(V)$), or to the unique point of $\bar{\mathbf{B}}$ belonging to $\Sigma(X^{\mathrm{an}}, V)$ (if x belongs to an open ball $\mathbf{B}$ in $\Delta(V)$). By a theorem of Berkovich, the map τ is in fact a strong deformation retraction. In particular, X^{an} and Γ have the same homotopy type (and the homotopy type of Γ is independent of the choice of V).

If V is a semistable vertex set for an affine curve X', there is a similarly defined retraction map $\tau : (X')^{\mathrm{an}} \twoheadrightarrow \Sigma((X')^{\mathrm{an}}, V)$.

4.5. **The genus formula.** For $x \in X^{\mathrm{an}}$ of type 2, recall that the residue field $\widetilde{\mathcal{H}(x)}$ of $\mathcal{H}(x)$ (which we like to call the "double residue field" of x) has transcendence degree one over the residue field κ of $\mathbb{K}$. (Recall that $\mathbb{K}$ is algebraically closed by hypothesis, and therefore κ is algebraically closed as well.) Let C_x be the unique smooth projective curve over κ with function field $\widetilde{\mathcal{H}(x)}$, and let $g_x = \omega(x)$ be the genus of C_x.

If V is any semistable vertex set for X^{an}, then $g_x = 0$ for all $x \notin V$ and the *genus formula* asserts that

$$g(X) = \sum_{x \in V} g_x + g(\Sigma(X^{\mathrm{an}}, V)).$$

4.6. **Tangent vectors.** There is a canonical metric on $\mathbf{H}(X^{\mathrm{an}}) := X^{\mathrm{an}} \setminus X(\mathbb{K})$ which restricts to the metric on $\Gamma = \Sigma(X^{\mathrm{an}}, V)$ for any semistable vertex set V for X^{an}. Note, however, that the metric topology on $\mathbf{H}(X^{\mathrm{an}})$ is much finer than the subspace topology inherited from X^{an}. A *geodesic segment* starting at $x \in \mathbf{H}(X^{\mathrm{an}})$ is an isometric embedding $\alpha : [0, a] \to \mathbf{H}(X^{\mathrm{an}})$ for some $a > 0$ such that $\alpha(0) = x$. We say that two geodesic segments starting at x are *equivalent* if they agree on a neighborhood of 0. A *tangent direction* at a point $x \in \mathbf{H}(X^{\mathrm{an}})$ is an equivalence class of geodesic segments starting at x. We denote by $T_x = T_x(X^{\mathrm{an}})$ the set of all tangent directions at x.

Every point $x \in X^{\mathrm{an}}$ has a neighborhood base consisting of simply connected open sets. For any simply connected neighborhood U of $x \in X^{\mathrm{an}}$, there is a natural bijection between T_x and the connected components of $U \setminus \{x\}$. Since $U \setminus \{x\}$ is connected for any type 1 point $x \in X(\mathbb{K}) \subset X^{\mathrm{an}}$, it is natural to say that there is only one tangent direction at x when x is of type 1. For x of type 4 we also have $|T_x| = 1$, and for x of type 3 we have $|T_x| = 2$.

If x is of type 2, on the other hand, then $|T_x|$ is infinite. In fact, there is a canonical bijection between T_x and $C_x(\kappa)$, the set of closed points on the κ-curve C_x corresponding to x. Equivalently, there is a canonical bijection between T_x and the set of discrete valuations on $\widetilde{\mathcal{H}(x)} = \kappa(C_x)$ which are trivial on κ. For $\vec{\nu} \in T_x$, we denote by $\mathrm{ord}_\nu : \kappa(C_x)^\times \to \mathbb{Z}$ the corresponding discrete valuation, so that for every nonzero rational function $\widetilde{f} \in \kappa(C_x)$ we have $\mathrm{ord}_\nu(\widetilde{f}) = \mathrm{ord}_{p_\nu}(\widetilde{f})$, where p_ν is the closed point of C_x corresponding to $\vec{\nu}$.

4.7. **The metrized complex associated to a semistable vertex set.** Given a semistable vertex set V for X^{an}, there is a canonical corresponding metrized complex $\mathfrak{C}V$ of κ-curves. Indeed, we have already defined a metric graph $\Gamma = \Sigma(X^{\mathrm{an}}, V)$ corresponding to V, and Γ comes equipped with a natural model G whose vertices are the points of V and whose edges correspond bijectively to the open annuli in the semistable decomposition associated to V. For each $v \in V$, we have also defined a smooth projective curve C_v over κ. It remains to specify, for each $v \in V$, a bijection ψ_v from the edges of G incident to v to a subset $\mathcal{A}_v$ of

$C_v(\kappa)$. Given such an edge e, we define $\psi(e)$ to be the point of $C_v(\kappa)$ corresponding to the tangent direction at v defined by e.

Remark 4.1. Passing to a larger semistable vertex set is compatible with linear equivalence of divisors and does not change the rank of divisors. More precisely, given two semistable vertex sets $V_1 \subseteq V_2$, there is a natural (retraction) map $\tau_{2,1} : \mathfrak{C}V_2 \to \mathfrak{C}V_1$ which induces a surjective map $\tau_{2,1} : \mathrm{Div}(\mathfrak{C}V_2) \to \mathrm{Div}(\mathfrak{C}V_1)$. In addition, $\tau_{2,1}$ sends principal divisors on $\mathfrak{C}V_2$ to principal divisors on $\mathfrak{C}V_1$ and so induces a map $\tau_{2,1} : \mathrm{Pic}(\mathfrak{C}V_2) \to \mathrm{Pic}(\mathfrak{C}V_1)$. It is easy to see that $\tau_{2,1}$ is an isomorphism and that for any divisor $\mathcal{D}$ on $\mathfrak{C}V_2$, one has $r_{\mathfrak{C}V_2}(\mathcal{D}) = r_{\mathfrak{C}V_1}(\tau_{2,1}(\mathcal{D}))$ (cf. Remark 2.2). As a consequence, there is a canonical group $\mathrm{Pic}\,\mathfrak{C}X^{\mathrm{an}}$ associated to X^{an}, defined as $\mathrm{Pic}(\mathfrak{C}V)$ for any semistable vertex set V, together with a canonical rank function $r : \mathrm{Pic}\,\mathfrak{C}X^{\mathrm{an}} \to \mathbb{Z}$.

4.8. **Semistable vertex sets and semistable models.** Semistable vertex sets for X^{an} correspond bijectively to semistable R-models $\mathfrak{X}$ for X. (Since we are assuming that X is proper, it does not matter if we work with algebraic or formal models for X; for simplicity, we will work with algebraic models.) In order to explain this in more detail, recall that a connected reduced algebraic curve over κ is called *semistable* if all of its singularities are ordinary double points, and is called *strongly semistable* if in addition its irreducible components are all smooth. A *(strongly) semistable model* for X is a flat and integral proper relative curve $\mathfrak{X}$ over R whose generic fiber is isomorphic to X and whose special fiber $\bar{\mathfrak{X}}$ is a (strongly) semistable curve.

Given a semistable model $\mathfrak{X}$ for X, there is a canonical associated reduction map $\mathrm{red} : X(\mathbb{K}) \to \bar{\mathfrak{X}}(\kappa)$ which is defined using the natural bijection between $X(\mathbb{K})$ and $\mathfrak{X}(R)$. This extends naturally to a map $\mathrm{red} : X^{\mathrm{an}} \to \bar{\mathfrak{X}}$. Let $\bar{\mathfrak{X}}^{\mathrm{gen}}$ be the set of generic points of irreducible components of $\bar{\mathfrak{X}}$. Then $V(\mathfrak{X}) := \mathrm{red}^{-1}(\bar{\mathfrak{X}}^{\mathrm{gen}})$ is a finite set of type 2 points mapping bijectively onto $\bar{\mathfrak{X}}^{\mathrm{gen}}$, and moreover $V(\mathfrak{X})$ is a semistable vertex set for X^{an}. The association $\mathfrak{X} \mapsto V(\mathfrak{X})$ gives a *bijection* between semistable formal models of X and semistable vertex sets for X^{an}. The model $\mathfrak{X}$ is strongly semistable if and only if $V(\mathfrak{X})$ is a strongly semistable vertex set.

4.9. **The metrized complex associated to a semistable model.** Since there is a natural bijection between (strongly) semistable vertex sets for X^{an} and semistable models for X, it follows that we can canonically associate to any semistable model $\mathfrak{X}$ a corresponding metrized complex of κ-curves $\mathfrak{C}\mathfrak{X}$. However, one can give a more direct description of $\mathfrak{C}\mathfrak{X}$ as follows.

Let G be the dual graph of $\bar{\mathfrak{X}}$, so that vertices of G correspond to irreducible components of $\bar{\mathfrak{X}}$ and edges of G correspond to intersections between irreducible components. If x^e is the ordinary double point of $\bar{\mathfrak{X}}$ corresponding to an edge e of G, the *formal fiber* $\mathrm{red}^{-1}(x^e)$ is isomorphic to an open annulus $\mathbf{A}$. We define the *length of the edge* e to be the length of the skeleton of $\mathbf{A}$, i.e., the *modulus* $\log(R) - \log(r)$ of $\mathbf{A} \cong \{x \in (\mathbf{A}^1)^{\mathrm{an}} \mid r < |T|_x < R\}$. (The modulus of an open annulus is well-defined independent of the choice of such an analytic isomorphism.) In this way, we have defined a metric graph $\Gamma = \Gamma_{\mathfrak{X}}$ associated to $\mathfrak{X}$ together with a model G. The irreducible components C_v of $\bar{\mathfrak{X}}$ correspond bijectively to the vertices v of G, and we let $A_v \subset C_v$ be the finite set of double points of $\bar{\mathfrak{X}}$ contained in C_v, so that there is a natural bijection between A_v and the edges of G incident to v. In this way we have defined a metrized complex $\mathfrak{C}\mathfrak{X}$ canonically associated to $\mathfrak{X}$. The metrized complex $\mathfrak{C}\mathfrak{X}$ is the same as the metrized complex $\mathfrak{C}V$ which we previously associated to $V = V(\mathfrak{X})$.

One can show that essentially every metrized complex of curves comes from this construction. The following result is proved in [ABBR]:

Theorem 4.2. *Let $\mathfrak{C}$ be a metrized complex of κ-curves whose edge lengths are contained in the value group of $\mathbb{K}$. There exists a smooth, proper, connected curve X over $\mathbb{K}$ and a semistable model $\mathfrak{X}$ for X such that $\mathfrak{C} \cong \mathfrak{C}\mathfrak{X}$.*

4.10. **Specialization of divisors from curves to metrized complexes.** As discussed earlier, there is a canonical embedding of $\Gamma_{\mathfrak{X}}$ in the Berkovich analytification X^{an} of X, as well as a canonical retraction map $\tau : X^{\mathrm{an}} \to \Gamma_{\mathfrak{X}}$. In addition, there is a canonical reduction map $\mathrm{red} : X^{\mathrm{an}} \to \bar{\mathfrak{X}}$ sending $X(\mathbb{K})$ to surjectively onto the closed points of $\bar{\mathfrak{X}}$. The retraction map τ induces by linearity a specialization map $\tau_* : \mathrm{Div}(X) \to \mathrm{Div}(\Gamma_{\mathfrak{X}})$ which is studied in [B]. We can promote this to a map $\tau_*^{\mathfrak{C}\mathfrak{X}}$ whose target is the larger group $\mathrm{Div}(\mathfrak{C}\mathfrak{X})$ as follows. If U is a subset of X^{an} and $D = \sum a_x(x) \in \mathrm{Div}(X)$, we define $U \cap D := \sum_{x \in U} a_x(x)$. For $v \in V$, define $\sigma_v : \mathrm{Div}(X) \to \mathrm{Div}(C_v)$ by

$$\sigma_v(D) = \mathrm{red}\left(\mathrm{red}^{-1}(C_v(\kappa) \setminus \mathcal{A}_v) \cap D\right)$$

and put $\sigma(D) = \sum_{v \in V} \sigma_v(D) \in \mathrm{Div}(\mathfrak{C}\mathfrak{X})$. Then

$$\tau_*^{\mathfrak{C}\mathfrak{X}}(D) := \tau_*(D) \oplus \sigma(D).$$

Equivalently, one can describe $\tau_*^{\mathfrak{C}\mathfrak{X}}$ as follows. If $P \in X^{\mathrm{an}}$ satisfies $\tau(P) = v \in V$ then either P is the unique point of X^{an} with $\mathrm{red}(P)$ equal to the generic point of C_v, or else $\mathrm{red}(P)$ is a nonsingular closed point of C_v. In the latter case we let $\bar{P} \in C_v(\kappa) \setminus \mathcal{A}_v$ denote the reduction of P. The map $\tau_*^{\mathfrak{C}\mathfrak{X}} : \mathrm{Div}(X) \to \mathrm{Div}(\mathfrak{C}\mathfrak{X})$ is obtained by linearly extending the map $\tau^{\mathfrak{C}\mathfrak{X}} : X(\mathbb{K}) \to \mathrm{Div}(\mathfrak{C}\mathcal{X})$ defined by

$$\tau^{\mathfrak{C}\mathfrak{X}}(P) = \begin{cases} \tau(P) & \tau(P) \notin V \\ \tau(P) \oplus \bar{P} & \tau(P) \in V. \end{cases}$$

4.11. **Reduction of rational functions and Poincaré-Lelong.** Let $x \in X^{\mathrm{an}}$ be a point of type 2. Given a nonzero rational function f on X, choose $c \in \mathbb{K}^\times$ such that $|f(x)| = |c|$. Define $f_x \in \kappa(C_x)^\times$ to be the image of $c^{-1}f$ in $\widetilde{\mathcal{H}(x)} \cong \kappa(C_x)$. Although f_x is only well-defined up to multiplication by an element of $\kappa^\times$, its divisor $\mathrm{div}(f_x)$ is canonical and the resulting map $\mathrm{Prin}(X) \to \mathrm{Prin}(C_x)$ is a homomorphism. Note also that if ord is a discrete valuation on $\kappa(C_x)$ which is trivial on κ then $\mathrm{ord}(f_x)$ is intrinsic to f.

If H is a $\mathbb{K}$-linear subspace of $\mathbb{K}(X)$, the collection of all possible reductions of nonzero elements of H, together with $\{0\}$, forms a κ-vector space H_x. For later use we note the following elementary lemma, which says that $\dim H = \dim H_x$:

Lemma 4.3. *Let X be a smooth proper curve over $\mathbb{K}$, and $x \in X^{\mathrm{an}}$ a point of type 2. The κ-vector space H_x defined by the reduction to $\widetilde{\mathcal{H}(x)}$ of an $(r+1)$-dimensional $\mathbb{K}$-subspace $H \subset \mathbb{K}(X)$ has dimension $r+1$.*

Proof. The inequality $\dim_\kappa(H_x) \leq r+1$ follows from the observation that if $f_0, \ldots, f_n$ are linearly dependent in H, with $|f_i| = 1$ for all i, then the reductions $f_{0,x}, \ldots, f_{n,x}$ are linearly dependent in H_x. To prove $\dim_\kappa(H_x) \geq r+1$, choose a $\mathbb{K}$-basis $f_0, \ldots, f_r$ for H consisting of elements of norm 1 and let $f_{0,x}, \ldots, f_{n,x}$ in H_x be the corresponding reductions. It is not true that $f_{0,x}, \ldots, f_{n,x}$ necessarily form a basis of H_x: for example, if f and g are linearly independent functions of norm 1 and c is any scalar with $|c| < 1$, then f and $f + cg$ are linearly independent but they have the same reduction. However, by the following procedure one can construct another basis $g_0, \ldots, g_r$ of H with $|g_i| = 1$ for all i such that the reductions $g_{i,x}$ are linearly independent over κ. Define $g_0 = f_0$, and recursively define g_{i+1} as follows: among all the rational functions of the form $f_{i+1} - \sum_{j \leq i} a_j g_j$ with $|a_j| \leq 1$ for all j, choose one, call it h_{i+1}, whose norm is minimal. Note that $|f_{i+1} - \sum_{j \leq i} a_j g_j| \leq 1$ so such an element exists by compactness. Let c_{i+1} be a scalar with $|c_{i+1}| = |h_{i+1}|$ and define $g_{i+1} = c_{i+1}^{-1} h_{i+1}$. It is clear that the $\{g_i\}$ form a basis of H. We claim that the reductions $g_{i,x}$ form a basis of H_x. Indeed, if not then there exists a minimal index i and scalars b_j with $|b_j| \leq 1$ for $j = 0, \ldots, i$ such that $|g_{i+1} - \sum_{j \leq i} b_j g_j| < 1$. But this implies that $|f_{i+1} - \sum_{j \leq i}(a_j + c_{i+1} b_j) g_j| < c_{i+1}$, contradicting the choice of h_{i+1}. $\square$

We say that a function $F : X^{\mathrm{an}} \to \mathbb{R}$ is *piecewise linear* if for any geodesic segment $\alpha : [a,b] \hookrightarrow \mathbf{H}(X^{\mathrm{an}})$, the pullback map $F \circ \alpha : [a,b] \to \mathbb{R}$ is piecewise linear. The *outgoing slope* of a piecewise linear function F at a point $x \in \mathbf{H}(X^{\mathrm{an}})$ along a tangent direction $\vec{v} \in T_x$ is defined to be

$$d_{\vec{v}} F(x) = \lim_{\epsilon \to 0} (F \circ \alpha)'(\epsilon),$$

where $\alpha : [0,a] \hookrightarrow X^{\mathrm{an}}$ is a geodesic segment starting at x which represents $\vec{\nu}$. A piecewise linear function F is called *harmonic* at a point $x \in \mathbf{H}(X^{\mathrm{an}})$ if the outgoing slope $d_{\vec{\nu}} F(x)$ is zero for all but finitely many $\vec{\nu} \in T_x$, and $\sum_{\vec{\nu} \in T_x} d_{\vec{\nu}} F(x) = 0$.

Theorem 4.4 (Non-Archimedean Poincaré-Lelong Theorem). *Let f be a nonzero rational function on X, let X' be an open affine subset of X on which f has no zeros or poles, and let $F = -\log|f| : (X')^{\mathrm{an}} \to \mathbb{R}$. Let V be a semistable vertex set of X' and let $\Sigma = \Sigma(X', V)$ be its skeleton. Then:*

1. *$F = F \circ \tau_\Sigma$ where $\tau_\Sigma : (X')^{\mathrm{an}} \to \Sigma$ is the retraction.*
2. *F is piecewise linear with integer slopes, and F is linear on each edge of Σ.*
3. *If x is a type-2 point of X^{an} and $\vec{\nu} \in T_x$, then $d_{\vec{\nu}} F(x) = \mathrm{ord}_v(f_x)$.*
4. *F is harmonic at all $x \in \mathbf{H}(X^{\mathrm{an}})$.*
5. *Let $x \in X \setminus X'$, let e be the ray in Σ whose closure in X^{an} contains x, let $y \in V$ be the other endpoint of e, and let $\vec{\nu} \in T_y$ be the tangent direction represented by e. Then $d_{\vec{\nu}} F(y) = \mathrm{ord}_x(f)$.*

To each nonzero rational function f on X and each semistable vertex set V for X, there is a corresponding rational function $\mathfrak{f} = (f_\Gamma, \{f_v\}_{v \in V})$ on $\mathfrak{CX}$, where f_Γ is the restriction to $\Gamma = \Sigma(X^{\mathrm{an}}, V)$ of the piecewise linear function $F = \log|f|$ on X^{an}. (We take $F = \log|f|$ instead of $-\log|f|$ because we have defined $\mathrm{ord}_F(u)$ to be the sum of the outgoing slopes of F at u. One could equally well take the opposite convention, defining $\mathrm{ord}_F(u)$ to be *minus* the sum of the outgoing slopes of F at u, and then defining F to be $-\log|f|$. Such a modification would also necessitate a change of sign in our definition of $\mathrm{div}_v(F)$.)

As an application of Theorem 4.4, we obtain the following important formula:

Theorem 4.5. *For every nonzero rational function f on X,*

$$\tau_*^{\mathfrak{CX}}(\mathrm{div}(f)) = \mathrm{div}(\mathfrak{f}).$$

Proof. Consider a nonzero rational function f on X and let D be the support of $\mathrm{div}(f)$ in X. We apply Theorem 4.4 to $X' = X \setminus D$, the invertible rational function f on X', and a semistable vertex V' of X' which contains V. Let $\Sigma = \Sigma(X'^{\mathrm{an}}, V')$ and $\Gamma = \Sigma(X, V)$. Since V is a semistable vertex set for X, the closure of $\Sigma \setminus \Gamma$ in Σ is a disjoint union of metric trees $\mathcal{T}_1, \ldots, \mathcal{T}_s$. For each tree $\mathcal{T}_i$, denote by $x_1^i, \ldots, x_{n_i}^i$ all the points in $D \subset X(\mathbb{K})$ which are in the closure of rays in $\mathcal{T}_i$, and let $\overline{\mathcal{T}_i} = \mathcal{T}_i \cup \{x^i_1, \dots, x^i_{n_i}\}$. Denote by y^i the unique point of $\Gamma \cap \overline{\mathcal{T}_i}$; y_i is the reduction in Γ of all the points $x^i_{n_i}$. Note that the points y_i are all different. In addition, if $y^i \in V$, denote by z^i_j the reduction of x^i_j in C_{y_i}, so that $\tau_*^{\mathfrak{C}\mathfrak{X}}(x^i_j) = y^i \oplus z^i_j$.

For each $i \in \{1, \dots, s\}$, the restriction of $F = \log|f|$ to $\mathcal{T}_i$ is harmonic at every point of $\overline{\mathcal{T}_i} \setminus \{y_i, x^i_1, \dots, x^i_{n_i}\}$ by (4). By (5), $\log|f|$ is linear in a sufficiently small neighborhood of x^i_j in $\mathcal{T}_i$, and has (constant) slope $\mathrm{ord}_{x^i_j}(f)$ along the unique tangent direction in this interval toward the point y_i. We infer that the order of $F|_{\mathcal{T}_i}$ at y^i is equal to $-\sum_{1 \le j \le n_i} \mathrm{ord}_{x^i_j}(f)$. In addition, by (1) F is constant on any connected component of $X^{\mathrm{an}} \setminus \Sigma$. Since F is harmonic at point y^i in Σ, we infer that, for the restriction f_Γ of F to Γ, we have $\mathrm{ord}_{y^i}(f_\Gamma) = -\mathrm{ord}_{y_i}(F|_{\mathcal{T}_i}) = \sum_{1 \le j \le n_i} \mathrm{ord}_{x^i_j}(f)$. Since f_Γ is linear on each edge of Γ, it follows that $\tau_*(\mathrm{div}(f)) = \mathrm{div}(f_\Gamma)$.

Suppose now that y_i is a point of V. For any tangent direction $\vec{\nu} \in T_{y_i}$, by (3) we have $d_{\vec{\nu}}F = -\mathrm{ord}_\nu(f_{y_i})$. Since F is locally constant on $X^{\mathrm{an}} \setminus \Sigma$, we infer that $d_{\vec{\nu}}F = 0$ for any $\vec{\nu} \in T_x$ which does not lie in Σ. This shows that $\mathrm{div}(f_{y_i})$ is supported on $\{z^i_1, \dots, z^i_{n_i}\} \cup \mathcal{A}_{y_i}$. It remains to prove the equality of coefficients. By (3) for any point ν in $\mathcal{A}_{y_i}$, we have $d_{\vec{\nu}} f_\Gamma + \mathrm{ord}_\nu(f_{y_i}) = 0$. This shows that the coefficient of ν in $\mathrm{div}(\mathfrak{f})$ is zero. For a point $\nu \in \{z^i_1, \dots, z^i_{n_i}\}$, let $\overline{\mathcal{T}}_{i,\nu}$ be the closure in Σ of the connected component of $\Sigma \setminus y_i$ which contains all the points x^i_j with $z^i_j = \nu$. By (1) and (4), F is harmonic at any point of $\overline{\mathcal{T}}_{i,\nu} \setminus (D \cup \{y_i\})$, and is linear of slope $\mathrm{ord}_{x^i_j}(f)$ in a sufficiently small neighborhood of x^i_j in $\mathcal{T}_{i,\nu}$ along the unique tangent direction in this interval toward the point y_i. By (3), we infer that along tangent direction $\vec{\nu}$ in T_{y_i} corresponding to ν we have $\mathrm{ord}_\nu(f_{y_i}) = -d_{\vec{\nu}}(F) = \sum_{x^i_j :\, z^i_j = \nu} \mathrm{ord}_{x^i_j}(f)$. We conclude that $\tau_*^{\mathfrak{C}\mathfrak{X}} \mathrm{div}(f) = \mathrm{div}(\mathfrak{f})$. ☐

In particular, it follows that $\tau_*^{\mathfrak{C}\mathfrak{X}}(\mathrm{Prin}(X)) \subseteq \mathrm{Prin}(\mathfrak{C}\mathfrak{X})$.

Remark 4.6. With the notation introduced in Remark 4.1, the map $\tau_*^{\mathfrak{C}\mathfrak{X}}$ induces a canonical homomorphism $\mathrm{Pic}(X) \to \mathrm{Pic}(\mathfrak{C}X^{\mathrm{an}})$ which is deserving of further study.

4.12. **Higher-dimensional metrized complexes.** We restrict our attention in this paper to metrized complexes of curves. Nevertheless, it is clear that many of our basic definitions and constructions should have higher-dimensional analogues. We sketch here a possible approach which we feel merits further investigation.

Let $\mathbb{K}$ be a complete and algebraically closed non-trivially valued non-Archimedean field, R the valuation ring of $\mathbb{K}$, and κ its (algebraically closed) residue field. Let $\mathfrak{X}$ be a proper and flat strongly semistable R-scheme with smooth and irreducible generic fiber $X/\mathbb{K}$. (By *strongly semistable*, we mean that the special fiber $\bar{\mathfrak{X}}$ has smooth irreducible components with simple normal crossings.) The *metrized complex of κ-schemes associated to* $\mathfrak{X}$ consists of the following data:

- The simplicial complex $G_{\mathfrak{X}}$ (the "dual complex") associated to $\bar{\mathfrak{X}}$.
- The skeleton $\Gamma_{\mathfrak{X}}$ of $\mathfrak{X}$ in the sense of [Ber04], which is the geometric realization of $G_{\mathfrak{X}}$ together with its canonical integer affine structure. (Berkovich proves in *loc. cit.* that there is a canonical embedding of $\Gamma_{\mathfrak{X}}$ in X^{an} as well as a strong deformation retraction $\tau : X^{\mathrm{an}} \to \Gamma_{\mathfrak{X}}$.)
- For each vertex (0-dimensional cell) v of $G_{\mathfrak{X}}$, the corresponding κ-variety X_v.

- For each face F of $G_{\mathfrak{X}}$ corresponding to a subset $S \subseteq V$, the corresponding non-empty scheme $X_F = \cap_{v \in S} X_v$ together with its compatible family of embeddings in $X_{F'}$ for $F' \subseteq F$. (In the terminology of some algebraic geometers, $\{X_F\}_{F \in G_{\mathfrak{X}}}$ is the "cubical variety" associated to $\bar{\mathfrak{X}}$.)

When κ has characteristic zero, it is known that every smooth proper algebraic variety $X/\mathbb{K}$ has a strongly semistable R-model. Because this is unknown when κ has positive characteristic, Berkovich has introduced in [Ber99, Ber04] an intricate procedure for generalizing the skeleton $\Gamma_{\mathfrak{X}}$ and the corresponding deformation retraction τ to *pluristable* R-models, which allows one to utilize de Jong's theory of alterations instead of resolution of singularities. It would be interesting to study line bundles and intersection theory on metrized complexes of κ-schemes associated to semistable and/or pluristable models.

5. SPECIALIZATION INEQUALITY

Let $\mathbb{K}$ be a complete and algebraically closed non-Archimedean field with non-trivial absolute value $|.|$, R the valuation ring of $\mathbb{K}$, and κ its (algebraically closed) residue field. Let $\mathcal{G} = \mathrm{val}(\mathbb{K}^\times)$ be the value group of $\mathbb{K}$. Let X be a smooth, proper, connected curve over $\mathbb{K}$ and let X^{an} be the Berkovich analytic space associated to X. Fix a strongly semistable R-model $\mathfrak{X}$ for X and let $\mathfrak{C}\mathfrak{X}$ be the metrized complex associated to $\mathfrak{X}$.

The following is a variant of Lemma 2.8 and Corollary 2.11 from [B]:

Theorem 5.1 (Specialization Lemma). *For every divisor* $D \in \mathrm{Div}(X)$,

$$r_X(D) \le r_{\mathcal{C}\mathfrak{X}}(\tau_*^{\mathcal{C}\mathfrak{X}}(D)).$$

Proof. Let $\mathcal{D} := \tau_*^{\mathfrak{C}\mathfrak{X}}(D)$. We need to prove that if $r_X(D) \ge k$, then $r_{\mathfrak{C}\mathfrak{X},\mathcal{G}}(\mathcal{D}) \ge k$ as well. The base case $k = -1$ is obvious. Now suppose $k = 0$, so that $r_X(D) \ge 0$. Then there exists an effective divisor $E \in \mathrm{Div}(X)$ with $D - E \in \mathrm{Prin}(X)$. Since $\tau_*^{\mathfrak{C}\mathfrak{X}}$ is a homomorphism and takes principal (resp. effective) divisors on X to principal (resp. effective) divisors on $\mathfrak{C}\mathfrak{X}$, we have $\mathcal{D} = \tau_*^{\mathfrak{C}\mathfrak{X}}(D) \sim \tau_*^{\mathfrak{C}\mathfrak{X}}(E) \ge 0$, so that $r_{\mathfrak{C}\mathfrak{X},\mathcal{G}}(\mathcal{D}) \ge 0$ as well.

We may therefore assume that $k \ge 1$. For each $v \in V$, let $R_v \subset C_v(\kappa) \setminus \mathcal{A}_v$ be a subset of size $g_v + 1$ and denote by R the union of the R_v. By Theorem A.1 in Appendix A, R is a rank-determining set in $\mathfrak{C}$. Let $\mathcal{E}$ be an effective divisor of degree k with support in R, so that $\mathcal{E} = \sum_{i=1}^k v_i \oplus x_i$ with $v_i \in V(G)$ and $x_i \in R_{v_i}$. For each i, there exists a point $P_i \in X(\mathbb{K})$ such that $\tau_*^{\mathfrak{C}\mathfrak{X}}(P_i) = Q_i$, and so there is an effective divisor E of degree k on X with $\tau_*^{\mathfrak{C}\mathfrak{X}}(E) = \mathcal{E}$. By our assumption on the rank of D, $r_X(D - E) \ge 0$, and so $r_{\mathfrak{C}\mathfrak{X}}(\mathcal{D} - \mathcal{E}) \ge 0$. Since this is true for any effective divisor of degree k with support in R, and since R is rank-determining, we infer that $r_{\mathfrak{C}\mathfrak{X},\mathcal{G}}(\mathcal{D}) \ge k$ as desired. □

Remark 5.2. If X is defined over a complete but not necessarily algebraically closed non-Archimedean field $\mathbb{K}$, then $r_X(D) = r_{X_{\mathbb{C}_{\mathbb{K}}}}(D)$ for every $D \in \mathrm{Div}(X)$ and thus one still has an analogue of Lemma 5.1.

Remark 5.3. There are many examples where the inequality in Theorem 5.1 can be strict. For example, if $g_v = 0$ for all $v \in V$ then $r_{\mathcal{C}\mathfrak{X}}(\tau_*^{\mathcal{C}\mathfrak{X}}(D)) = r_\Gamma(\tau_*(D))$ (cf. Remark 2.2) and thus the examples of strict inequality from [B] apply. At the other extreme, if $C = \bar{\mathfrak{X}}$ is smooth (so that G is a point) then the specialization lemma becomes the well-known semicontinuity statement $h^0(D) \le h^0(\bar{D})$, where $\bar{D} \in \mathrm{Div}(C)$ is the reduction of D, and it is clear that such an inequality can be strict; for example, take $D = (P) - (Q)$ where X has genus at least 1 and $P, Q \in X(\mathbb{K})$ are distinct points with the same reduction in $C(\kappa)$.

5.1. **Refined versions of weighted specialization lemma.** In this section we show how the divisor theory and specialization lemma on metrized complexes provide a natural framework for obtaining strong forms of weighted specialization lemma, such as those given in [AC]. We start by recalling some basic terminology concerning divisors on weighted metric graphs.

Let (Γ,ω) be a weighted metric graph Γ of genus $g(\Gamma,\omega)=g(\Gamma)+\sum_{v\in V} g_v$ with a model $G=(V,E)$ containing all the points of Γ of positive weight. Denote by $\mathcal{W}$ the divisor $\mathcal{W}=\sum_v g_v(v)$ on Γ. Define a metric graph $\Gamma^{\#}$ by attaching g_v loops of arbitrary positive length at each point $v\in\Gamma$. Observe that the genus of $\Gamma^{\#}$ coincides with $g(\Gamma,\omega)$.

Any divisor D in Γ defines a divisor on $\Gamma^{\#}$ with the same support and coefficients (by viewing Γ as a subgraph of $\Gamma^{\#}$). By an abuse of notation, we will also denote this divisor by D. However, we distinguish the rank of D in Γ and in $\Gamma^{\#}$ by using the notation $r_\Gamma(D)$ for the former and $r_{\Gamma^{\#}}$ for the latter. For any divisor $D\in\mathrm{Div}(\Gamma)$, define the weighted rank $r^{\#}$ of D by $r^{\#}(D):=r_{\Gamma^{\#}}(D)$. In the next subsection, we give a precise formula for calculating $r^{\#}(D)$.

5.1.1. *The behavior of the rank function under connected sums.* Consider two (unweighted) metric graphs Γ_1 and Γ_2, and suppose that two distinguished points $v_1\in\Gamma_1$ and $v_2\in\Gamma_2$ are given. The *wedge* or *direct sum* of (Γ_1,v_1) and (Γ_2,v_2), denoted $\widetilde{\Gamma}=\Gamma_1\vee\Gamma_2$, is the metric graph obtained by identifying the points v_1 and v_2 in the disjoint union of Γ_1 and Γ_2. Denote by $v\in\widetilde{\Gamma}$ the image of v_1 and v_2 in $\widetilde{\Gamma}$. (By abuse of notation, we will use v to denote both v_1 in Γ_1 and v_2 in Γ_2.) We refer to $v\in\widetilde{\Gamma}$ as a *cut-vertex* and to $\widetilde{\Gamma}=\Gamma_1\vee\Gamma_2$ as the decomposition corresponding to the cut-vertex v. There is an addition map $\mathrm{Div}(\Gamma_1)\oplus\mathrm{Div}(\Gamma_2)\to\mathrm{Div}(\widetilde{\Gamma})$ which associates to any pair of divisors D_1 and D_2 in $\mathrm{Div}(\Gamma_1)$ and $\mathrm{Div}(\Gamma_2)$, the divisor, D_1+D_2 on $\widetilde{\Gamma}$ defined by pointwise addition of coefficients in D_1 and D_2.

Let g_1, g_2, and $\widetilde{g}=g_1+g_2$ be the genera of Γ_1,Γ_2, and $\widetilde{\Gamma}$, respectively, and let $r_1(.)$, $r_2(.)$ and $\widetilde{r}(.)=r_{\Gamma_1\vee\Gamma_2}(.)$ be the corresponding rank functions. The following proposition provides a precise description of $\widetilde{r}(D_1+D_2)$ in terms of the rank functions r_1 and r_2 in Γ_1 and Γ_2. In order to state it, we introduce a function $\eta=\eta_{v,D_2}:\mathbb{N}\cup\{0\}\to\mathbb{Z}$ defined as follows. For any non-negative integer k, denote by $\eta(k)$ the smallest integer n such that $r_2\big(D_2+n(v)\big)=k$. (By the Riemann-Roch theorem for Γ_2, η is well-defined.)

Proposition 5.4. *Let D_1 be a divisor in* $\mathrm{Div}(\Gamma_1)$. *Then*

$$\widetilde{r}(D_1+D_2)=\min_{k\in\mathbb{N}\cup\{0\}}\Big(k+r_1\big(D_1-\eta(k)\,(v)\big)\Big). \tag{1}$$

Proof. We prove the equality of the two sides of Equation (1) by showing that both the inequalities $\leq$ and $\geq$ hold.

($\geq$) Let r be the right-hand term in (1). Let E be an arbitrary effective divisor of degree r on $\widetilde{\Gamma}$ and write $E=E_1+E_2$ for two divisors E_1 and E_2 of degree $r-k$ and k with support in Γ_1 and Γ_2, respectively. We need to show that there exists a rational function f such that $\mathrm{div}(f)+D_1+D_2-E\geq 0$. By the definition of η, there exists a rational function f_2 with support in Γ_2 such that $\mathrm{div}(f_2)+\eta(k)(v)+D_2-E_2\geq 0$. Since $r_1\big(D_1-\eta(k)(v)\big)\geq r-k$, there exists a rational function f_1 with support in Γ_1 such that $\mathrm{div}(f_1)+D_1-\eta(k)(v)-E_2\geq 0$. For the rational function $f=f_1+f_2$ on $\widetilde{\Gamma}$, we thus have $\mathrm{div}(f)+D_1+D_2-E\geq 0$.

($\le$) It will be enough to show that for each $k \in \mathbb{N}$, the inequality $\widetilde{r}(D_1+D_2) \le k + r_1\big(D_1 - \eta(k)(v)\big)$ holds. Let $\widetilde{r} = \widetilde{r}(D_1+D_2)$. We can obviously restrict to the case $k \le \widetilde{r}$. For each effective divisor E_2 of degree k on Γ_2, let f_{E_2} be a rational function on Γ_2 with the property that the coefficients of $\mathrm{div}(f_{E_2}) + D_2 - E_2$ outside v are all non-negative, and such that in addition, the coefficient of v in $\mathrm{div}(f_{E_2}) + D_2 - E_2$ is maximized among all rational functions with this property. Denote by n_{E_2} the coefficient of v in $\mathrm{div}(f_{E_2}) + D_2 - E_2$ and let $\bar{E}_2$ be an effective divisor of degree k on Γ_2 which has the smallest coefficient $n_{\bar{E}_2} =: n$, i.e.,

$$n_{\bar{E}_2} = \min\{n_{E_2} \,|\, E_2 \ge 0 : \, \deg(E_2) = k\}.$$

By the choice of n, for any effective divisor E_2 of degree k there exists a rational function f_{E_2} such that $\mathrm{div}(f_{E_2}) + D_2 - E_2 - n(v) \ge 0$. We deduce from the definition of η that $\eta(k) \le -n$.

By the choice of f_{E_2}, we infer that for any rational function f_2 on Γ with the property that $\mathrm{div}(f_2)+D_2-\bar{E}_2$ has non-negative coefficients outside v, the coefficient of v in $\mathrm{div}(f_2)+D_2-\bar{E}_2$ is at most $-\eta(k)$. For any divisor E_1 of degree $\widetilde{r}-k$ on Γ_1, there exists a rational function f_1+f_2, with f_i supported on Γ_i, such that $\mathrm{div}(f_1) + \mathrm{div}(f_2) + D_1 + D_2 - E_1 - E_2 \ge 0$. Since the coefficient of v in $\mathrm{div}(f_2) + D_2 - \bar{E}_2$ is at most $-\eta(k)$, the coefficient of v in $\mathrm{div}(f_1) + D_1 - E_1$ must be at least $\eta(k)$, and all the other coefficients of $\mathrm{div}(f_1) + D_1 - E_1$ must be be non-negative, i.e., $\mathrm{div}(f_1) + D_1 - E_1 - \eta(k)(v) \ge 0$. Since this is the case for any $E_1 \ge 0$ of degree $\widetilde{r} - k$, we conclude that $r_1\big(D_1 - E_1 - \eta(k)(v)\big) \ge \widetilde{r} - k$. □

Corollary 5.5. *Let (Γ, ω) be a weighted metric graph and $\Gamma^{\#}$ the associated virtual graph. For every $D \in \mathrm{Div}(\Gamma)$, we have*

$$r^{\#}(D) = \min_{0 \le E \le \mathcal{W}} \big(\deg(E) + r_\Gamma(D - 2E)\big). \tag{2}$$

Proof. By induction, it is enough to consider the case $\deg(\mathcal{W}) = 1$, i.e., $w_v = 1$ for some point $v \in \Gamma$ and $w_u = 0$ for all $u \ne v$. The result now follows from Proposition 5.4 applied to $\Gamma_1 = \Gamma$, Γ_2 a circle of arbitrary length, $\widetilde{\Gamma} = \Gamma^{\#} = \Gamma_1 \vee \Gamma_2$, $D_1 = D$ and $D_2 = 0$. An easy calculation shows that $\eta(0) = 0$, $\eta(1) = 2$, and $\eta(k) = k+1$ for $k > 1$. Since $r_\Gamma(D - (k+1)(v)) \ge r_\Gamma(D - k(v)) - 1$ for any k, we have

$$\min_{k \in \mathbb{N}} \Big(k + r_1\big(D_1 - \eta(k)(v)\big) \Big) = \min\{r_\Gamma(D), 1 + r_\Gamma(D - 2\,v)\},$$

which is precisely Equation (2) in the case $\deg(\mathcal{W}) = 1$. □

5.1.2. *Weighted specialization lemma.* We first define an analogue of the η function introduced in the previous section for a metrized complex $\mathfrak{C}$, taking into account the pointed curves $(C_v, \mathcal{A}_v)$, and we then generalize Proposition 5.4 to this setting.

Consider a divisor $\mathcal{D} = D_\Gamma \oplus \sum_v D_v$ on $\mathfrak{C}$. For every vertex v of G, define the function $\eta_v : \mathbb{N} \cup \{0\} \to \mathbb{N} \cup \{0\}$ as follows. For any $k \ge 0$, $\eta_v(k)$ is the minimum integer $n \ge 0$ such that there exists a divisor $\widetilde{D}_v$ of degree $n - \deg(D_v)$ supported on the points of $\mathcal{A}_v \subset C_v(\kappa)$ such that $r_{C_v}(D_v + \widetilde{D}_v) = k$. For any effective divisor $E = \sum_{v \in V(G)} E(v)(v)$ on G, define

$$\eta(E) := \sum_{v \in V(G)} \eta_v(E(v))\,(v). \tag{3}$$

Note that the definition of η depends on the data of the divisors $D_v \in \mathrm{Div}(C_v)$.

Proposition 5.6. *Let $\mathfrak{C}$ be a metrized complex of algebraic curves over κ. For any divisor $\mathcal{D} = D_\Gamma \oplus \sum_v D_v$, let η be the integer function defined by* (3)*. Then*

$$r_{\mathfrak{C}}(\mathcal{D}) \ \leq \ \min_{E\geq 0} \Big(\deg(E) + r_\Gamma(D_\Gamma - \eta(E)) \Big),$$

where the minimum is over all effective divisors E on G.

Proof. Fix an effective divisor E supported on the vertices of G. We need to prove that

$$r_{\mathfrak{C}}(\mathcal{D}) \leq \deg(E) + r_\Gamma(D_\Gamma - \eta(E)).$$

For the sake of contradiction, suppose this is not the case, i.e., that $\deg(E)+r_\Gamma(D_\Gamma-\eta(E)) < r_{\mathfrak{C}}(\mathcal{D})$. Since $r_{\mathfrak{C}}(\mathcal{D}) - \deg(E) > r_\Gamma(D_\Gamma - \eta(E))$, the theory of rank-determining sets (see Appendix A) implies that there exists a divisor E^0 of degree $r_{\mathfrak{C}}(\mathcal{D}) - \deg(E)$ supported on $\Gamma \setminus V(G)$ such that $|D - E^0_\Gamma - \eta(E)| = \emptyset$.

Consider all the effective divisors on $\mathfrak{C}$ of the form $(E^0_\Gamma + E) \oplus \sum_{v\in V} E_v$ where each E_v is an effective divisor of degree $E(v)$ on C_v. Note that $(E^0 + E) \oplus \sum_v E_v$ has degree $r_{\mathfrak{C}}(\mathcal{D})$, so that by the definition of the rank function, for any choice of effective divisors E_v on C_v there exists a rational function $\mathfrak{f}$ on $\mathfrak{C}$ such that $\mathrm{div}(\mathfrak{f}) + (D_\Gamma - E^0 - E) \oplus \sum_v D_v - E_v \geq 0$. Denote by f_Γ and f_v the different parts of $\mathfrak{f}$. Since $\mathrm{div}(f_\Gamma) + D_\Gamma - E_0 - E \geq 0$, one can assume that for a generic choice of the divisors E_v, the slope of f_Γ along the edges incident to each vertex of G is fixed. This shows the existence of a divisor $\widetilde{D}_v$ supported on the points $\mathcal{A}_v \subset C_v(\kappa)$ such that $D_v + \widetilde{D}_v$ has rank at least $E(v) = \deg(E_v)$ on C_v. By the definition of η_v, $\widetilde{D}_v$ has degree at least $\eta_v(E(v)) - \deg(D_v)$. Since the degree of $\widetilde{D}_v$ is precisely the total sum of the slopes of f_Γ along the edges adjacent to v, which is by definition $\mathrm{ord}_v(f_\Gamma)$, for any $v \in V(G)$ we obtain the inequalities $\mathrm{ord}_v(f_\Gamma) \geq \eta_v(E(v)) - \deg(D_v)$, i.e., $\mathrm{ord}_v(f_\Gamma) + \deg(D_v) - \eta_v(E(v)) \geq 0$. In addition, for the other points u of Γ we must have $\mathrm{ord}_u(f_\Gamma) + D(u) - E^0_\Gamma(u) \geq 0$. This implies that $|D - E^0_\Gamma - \eta(E)| \neq \emptyset$, a contradiction. □

Let $\mathbb{K}$ be a complete and algebraically closed non-Archimedean field with non-trivial absolute value $|.|$, R the valuation ring of $\mathbb{K}$, and and κ its residue field. Let $\mathcal{G} = \mathrm{val}(\mathbb{K}^\times)$ be the value group of $\mathbb{K}$. Let X be a smooth, proper, and geometrically integral curve over $\mathbb{K}$ and fix a semistable R-model $\mathfrak{X}$. Let (Γ, ω) denote the skeleton of $\mathfrak{X}$ together with the natural weight function $\omega(x) = g_x$ (see Section 4), and let $\mathcal{W} = \sum_{v\in V} g_v(v)$.

Theorem 5.7 (Weighted specialization inequality)**.** *For every divisor $D \in \mathrm{Div}(X)$,*

$$r_X(D) \leq r^{\#}(\tau_*(D)).$$

In addition, if the marked curves $(C_v, \mathcal{A}_v)$ in $\mathfrak{C}\mathfrak{X}$ are general in their moduli, then for generic effective divisors D on X we have

$$r_X(D) \leq \min_{0\leq E\leq 1} \{\deg(E) + r_\Gamma(\tau_*(D) - \eta(E))\},$$

where $E \leq 1$ means all the coefficients are at most one. Moreover, for any vertex v, $\eta_v(0) = \min\{\tau_*(D)(v), g_v\}$, and $\eta_v(1) = g_v + 1$.
In particular, if for any v, $\tau_*(D)(v) \geq g_v$, then $r_X(D) \leq r_\Gamma(\tau_*(D) - \mathcal{W})$.

Proof. Let $\mathcal{D} = \tau_*^{\mathfrak{C}\mathfrak{X}}(D) = D_\Gamma \oplus \sum_v D_v$ and note that $D_\Gamma = \tau_*(D)$. Applying Proposition 5.6, we obtain

$$r_{\mathfrak{C}\mathfrak{X}}(\tau_*^{\mathfrak{C}\mathfrak{X}}(D)) \leq \min_{E\geq 0} \Big(\deg(E) + r_\Gamma\big(\tau_*(D) - \eta(E)\big) \Big). \tag{4}$$

Define the function $\bar{\eta}_v : \mathbb{N} \cup \{0\} \to \mathbb{N} \cup \{0\}$ by $\bar{\eta}_v(k) = 2k$ for $k \leq g_v$ and $\bar{\eta}_v(k) = k + g_v$ for $k \geq g_v$. By the definition of η_v and Clifford's theorem for the curve C_v, for any integer $k \geq 0$ we have $\eta_v(k) \geq \bar{\eta}_v(k)$. Combining these inequalities and (4), we get the inequality

$$r_{\mathcal{C}\mathfrak{X}}(\tau_*^{\mathcal{C}\mathfrak{X}}(D)) \leq \min_{E \geq 0} \Big(\deg(E) + r_\Gamma\big(\tau_*(D) - \bar{\eta}(E)\big) \Big). \tag{5}$$

Note that for $E(v) > g_v$, we have $\bar{\eta}_v(E(v)) = g_v + E(v)$ and so, whatever the other coefficients of E are, we have

$$\begin{aligned} \deg(E) + r_\Gamma\big(\tau_*(D) - \bar{\eta}(E)\big) &= \deg(E - (v)) + 1 + r_\Gamma\big(D - \bar{\eta}(E - (v)) - (v)\big) \\ &\geq \deg(E - (v)) + r_\Gamma(D - \bar{\eta}(E - (v))). \end{aligned}$$

In other words, in taking the minimum in (5), one can restrict to effective divisors E with $E(v) \leq g_v$ for all v, i.e., $0 \leq E \leq \mathcal{W}$. Finally, applying (5) and Lemma 5.1, we obtain the inequality

$$r_X(D) \leq r^{\#}(\tau_*(D)).$$

This proves the first part of the theorem.

To get the second statement, suppose that all the marked curves $(C_v, \mathcal{A}_v)$ are general in their moduli and that D is an effective divisor supported on general points of X. Then all the divisors D_v are also supported on general points of C_v. In this case, $\eta_v(k) = g_v + k$ for any non-negative integer $k \geq 1$ and any $v \in V(G)$ (this follows, for example, from the dimension part of the Brill-Noether theorem). In addition, we have $\eta_v(0) = \min\{\deg(D_v), g_v\}$. Thus, as in the proof of the first part, in dealing with the minimum in (4) we can restrict to $0 \leq E \leq 1$. In the case $\deg(D_v) \geq g_v$ for any v, by the same argument we get the stronger inequality $r_X(D) \leq r_\Gamma(\tau_*(D) - \mathcal{W})$. $\square$

5.2. **Some applications.** We give some direct applications of the results of this section.

5.2.1. *Specialization of canonical divisors.* Let K_X be a canonical divisor on the curve $X/\mathbb{K}$. By Theorem 5.1, we have $r_{\mathfrak{C}\mathfrak{X}}\left(\tau_*^{\mathfrak{C}\mathfrak{X}}(K_X)\right) \geq r(K_X) = g(X) - 1 = g(\mathfrak{C}\mathfrak{X}) - 1$. By the Riemann-Roch theorem for metrized complexes (Theorem 3.2), the divisor $\tau_*^{\mathfrak{C}\mathfrak{X}}(K_X) - \mathcal{K}$ has degree zero and non-negative rank, and thus $\tau_*^{\mathfrak{C}\mathfrak{X}}(K_X) \sim \mathcal{K}$. In particular, we have $\tau_*(K_X) \sim K^{\#}$.

Remark 5.8. For discretely valued R, the fact that $\tau_*(K_X)$ is linearly equivalent to $K^{\#}$ was proved in [B, Lemma 4.19] by a completely different argument based on the adjunction formula for arithmetic surfaces. It is not obvious how to extend such an intersection-theoretic argument to the general (not necessarily discretely valued) case treated above.

5.2.2. *Brill-Noether theory.* The following result is an immediate consequence of Theorem 4.2 combined with Theorem 5.1.

Corollary 5.9 (Brill-Noether existence theorem). *Let κ be a field and let g, r, d be nonnegative integers. If the Brill-Noether number $\rho_d^r(g) := g - (r+1)(g-d+r)$ is nonnegative, then for every metrized complex $\mathfrak{C}$ of κ-curves with $g(\mathfrak{C}) = g$, there exists a divisor $\mathcal{D} \in \mathrm{Div}_+^d(\mathfrak{C})$ such that $r_{\mathfrak{C}}(\mathcal{D}) \geq r$.*

Remark 5.10. In particular, combining the original Specialization Lemma from [B] with Theorem 4.2 yields a new proof of [B, Theorem 3.20], which says that if $\rho_d^r \geq 0$ then for every metrized graph of genus g there exists $D \in \mathrm{Div}_+^d(\Gamma)$ such that $r(D) \geq r$. The present proof is more direct in that one does not need the approximation arguments given in Lemma 3.17 and Corollary 3.18 of *loc. cit.*

Remark 5.11. The notion of Brill-Noether rank on metric graphs introduced in [LPP] and the corresponding specialization lemma extend to the context of metrized complexes as well. Let $\mathfrak{C}$ be a metrized complex of κ-curves, and let $W_d^r(\mathfrak{C}) \subset \mathrm{Pic}^d(\mathfrak{C})$ be the subset of all divisors of degree d and rank at least r. Define $w_d^r(\mathfrak{C})$ to be the largest integer n such that for every effective divisor $\mathcal{E}$ of degree $r+n$ on $\mathfrak{C}$, there exists $\mathcal{D} \in W_d^r(\mathfrak{C})$ such that $\mathcal{D} - \mathcal{E} \geq 0$. If $W_d^r(\mathfrak{C}) = \emptyset$, define $w_d^r(\mathfrak{C}) = -1$. If $\mathfrak{X}$ is a semistable R-model of a proper smooth curve $X/\mathbb{K}$ such that $\mathfrak{C}\mathfrak{X} = \mathfrak{C}$, and $W_d^r(X) \subset \mathrm{Pic}^d(X)$ denotes the Brill-Noether locus of X, then we have the inequality $w_d^r(\mathfrak{C}) \geq \dim W_d^r(X)$.

5.2.3. *Weierstrass points on metrized complexes of κ-curves.* Let $\mathfrak{C}$ be a metrized complex of κ-curves. By Remark 2.1, any point $x \in \mathfrak{C}$ defines an effective divisor of degree one on $\mathfrak{C}$ which we denote by $(x) \in \mathrm{Div}_+^1(\mathfrak{C})$. We say that a point x in $\mathfrak{C}$ is a *Weierstrass point* if $r_{\mathfrak{C}}\big(g(\mathfrak{C})(x)\big) \geq 1$. By Theorem 5.1, any metrized complex of κ-curves of genus at least two has a Weierstrass point.

6. Limit linear series

A powerful technique in algebraic geometry for dealing with problems about algebraic curves is degeneration to singular curves. Given the nice properties of the moduli space of stable curves and the stable reduction theorem [DM], it is particularly natural to consider degeneration of smooth curves to singular (semi-)stable curves. In addition, many interesting problems about algebraic curves can be formulated as problems about linear series. In this regard, it is also natural to consider the limit behavior of linear series in degenerating families of curves. This has been very successful in practice and has led to important advances in the study of algebraic curves (see [HM98] for an overview); we mention for example Brill-Noether theory [EH, Bi], the geometry of moduli space of curves [HM82, EH2], and Weierstrass points and their monodromy [EH3, EM, EH4].

Our aim in this section is to compare our divisor theory on metrized complexes with the theory of limit linear series for curves of compact type, introduced by Eisenbud-Harris [EH] as a way to handle degeneration of linear series on smooth curves to certain semistable curves, called curves of compact type. As we will see, our comparison results lead naturally to a generalization of the notion of limit linear series to nodal curves which are not necessarily of compact type.

6.1. Degeneration of linear series.

We start by reviewing some basic facts and definitions concerning the degeneration of linear series in families. Let Y be a smooth projective curve over a field κ. Recall that a linear series L of (projective) dimension r and degree d, or simply a $\mathfrak{g}_d^r$, over Y consists of a pair $(\mathcal{L}, H)$ with $\mathcal{L}$ an invertible sheaf of degree d on Y and H a subspace of $H^0(Y, \mathcal{L})$ of κ-dimension $r+1$.

Let $\phi : \mathcal{X} \to B = \mathrm{Spec}\, R$ be a regular smoothing of a strongly semistable curve X_0 over a discrete valuation ring R. This means that ϕ is proper, $\mathcal{X}$ is regular, the generic fiber $\mathcal{X}_\eta$ of ϕ is smooth, and the special fiber of ϕ is X_0. Denote by $G = (V, E)$ the dual graph of X_0 and let $\{X_v\}_{v \in V(G)}$ be the set of irreducible components of X_0.

Proposition 6.1. *Any line bundle $\mathcal{L}_\eta$ on $\mathcal{X}_\eta$ extends to $\mathcal{X}$. The restriction of any two extensions $\mathcal{L}_1$ and $\mathcal{L}_2$ of $\mathcal{L}_\eta$ to the special fiber X_0 are combinatorially equivalent (c.f. Section 2.2).*

Proof. The first part is a well-known consequence of the regularity of $\mathcal{X}$. For the second part, note that any two extensions $\mathcal{L}_1$ and $\mathcal{L}_2$ of $\mathcal{L}_\eta$ differ by a divisor supported on the special fiber X_0. In other words, there exist integers n_v for each vertex v of the dual graph of X_0 such that $\mathcal{L}_1 \simeq \mathcal{L}_2(\sum_v n_v X_v)$. It is now easy to see that, in the notation of Section 2.2, $\pi(\mathcal{L}_1|_{X_0})$ and $\pi(\mathcal{L}_2|_{X_0})$ differ by the divisor of the function $f_G : V(G) \to \mathbb{Z}$ which sends v to n_v. □

Consider now a $\mathfrak{g}^r_d$ $L_\eta = (\mathcal{L}_\eta, H_\eta)$ on $\mathcal{X}_\eta$ and choose an extension $\mathcal{L}$ of $\mathcal{L}_\eta$ to $\mathcal{X}$. Every section $s \in H_\eta$ uniquely extends to a global section $\bar{s}$ of $\mathcal{L}$, and the set of all such $\bar{s}$ forms a free R-submodule H of $H^0(\mathcal{X}, \mathcal{L})$. The restriction H_0 of H to the special fiber X_0 is an $(r+1)$-dimensional subspace of $H^0(X_0, \mathcal{L}_0)$. This shows the existence of the limit of a $\mathfrak{g}^r_d$ once the extension $\mathcal{L}$ of $\mathcal{L}_\eta$ is fixed. However, since the extension is not unique, the limit pair $(\mathcal{L}_0, H_0)$ is not unique. If one were to restrict to a single such extension, there would inevitably be a loss of important information; this can be seen, for example, in the study of limits of ramification points and in the study of limits of smooth hyperelliptic, and more generally, smooth d-gonal curves, see for example [HM82, EH, EM, R]. Some of the difficulties in formulating a good notion of limits for linear series are already manifest in the atypical behavior of linear series on reducible curves, since (as mentioned earlier) they do not satisfy many of the classical theorems which govern the behavior of (global sections of) linear series on irreducible curves.

The Eisenbud-Harris theory of limit linear series on curves of compact type [EH] provides a way to keep track of the geometry in the limit by choosing an *aspect* of the limit for each irreducible component, with some extra conditions relating these different aspects at nodes (see below for a precise definition). We refer to [O1, O2] for a relatively recent refinement of the Eisenbud-Harris theory, and to [E, EM] for attempts to generalize this theory to certain curves not of compact type.

6.2. **Limit linear series for curves of compact type.** We recall the definition of limit linear series for curves of compact type, following [EH].

We first need to recall the sequence of orders of vanishing of a $\mathfrak{g}^r_d$ on a smooth projective curve at a given point. Let Y be a smooth projective curve over an algebraically closed field κ and let $L = (\mathcal{L}, H)$ be a $\mathfrak{g}^r_d$ over Y. For any closed point p on Y and any section $f \in H$, denote by $\mathrm{ord}_p(f)$ the order of vanishing of f at p. The orders of vanishing at p of all the sections of $\mathcal{L}$ in H define a sequence of integers $0 \leq a^L_0(p) < a^L_1(p) < \cdots < a^L_{r-1}(p) < a^L_r(p)$. (This sequence is obtained by induction: let f_r be the section in H with the highest order of vanishing, and define f_{r-i} inductively as a section in H with the highest possible order of vanishing at p among all sections which are linearly independent of the first i sections $f_r, \ldots, f_{r-i+1}$. Then $a_i(p) = \mathrm{ord}_p(f_i)$.)

Let X_0 be a reduced strongly semistable curve over κ, let $G = (V, E)$ be the dual graph of X_0, and for any $v \in V$, let X_v be the corresponding irreducible component of X_0. When G is a tree, the curve X_0 is called of *compact type*: in this case, $\mathrm{Pic}(X_0)$ is compact.

(Recall that in general, we have an exact sequence of algebraic groups

$$0 \to H^1(G,\mathbb{Z}) \otimes \mathbb{G}_m \to \mathrm{Pic}^0(X_0) \to \prod_{v \in V_0} \mathrm{Pic}^0(X_v) \to 0.)$$

Let X_0 be a curve of compact type. A *crude limit* $\mathfrak{g}^r_d$ L over X_0 is by definition the data of a $\mathfrak{g}^r_d$ $L_v = (\mathcal{L}_v, H_v)$ over X_v for each vertex $v \in V$, called the X_v*-aspect* of L, such that the following property holds: If the two components X_u and X_v of X_0 meet at a node p corresponding to an edge $\{u,v\}$ in E, then

$$a_i^{L_v}(p) + a_{r-i}^{L_u}(p) \geq d \tag{6}$$

for all $0 \leq i \leq r$. A crude limit linear series is a *refined limit linear series* if all the inequalities in (6) are equalities.

Let Γ be the metric graph associated to G and let $\mathfrak{C}X_0$ be the regularization of X_0 (c.f. Section 2.2). For each vertex v of G, let $\overline{D}_v$ be a divisor of degree d on X_v in the divisor class defining $\mathcal{L}_v$ (i.e., $\mathcal{L}(\overline{D}_v) \simeq \mathcal{L}_v$). The subspace $H_v \subseteq H^0(X_v, \mathcal{L}(\overline{D}_v))$ can be naturally regarded as a subspace of $\kappa(X_v)$. Let $\mathfrak{r}$ be a fixed vertex of G and let $G_{\mathfrak{r}}$ denote G considered as a rooted tree with root $\mathfrak{r}$. For any vertex $u \neq \mathfrak{r}$ in G, consider the unique path from u to $\mathfrak{r}$ in G and let e_u be the adjacent edge to u on this path. Let x_u be the node of X_u corresponding to the edge e_u. Consider the divisor $\mathcal{D} = D_\Gamma \oplus \sum_{v \in V(G)} D_v$ of degree d on $\mathfrak{C}X_0$ defined by

$$D_\Gamma := d(\mathfrak{r}), \;\; D_{\mathfrak{r}} := \overline{D}_{\mathfrak{r}}, \;\; \text{and} \;\; D_u := \overline{D}_u - d(x_u) \text{ for } u \neq \mathfrak{r}. \tag{7}$$

We will see below that $\mathcal{D}$ has rank at least r in $\mathfrak{C}X_0$. In order to derive a more precise characterization of crude limits in terms of a suitable rank function, we need to introduce a refined notion of rank for divisors on a metrized complex of curves.

6.2.1. *A refined notion of rank for divisors on a metrized complex.* Let $\mathfrak{C}$ be a metrized complex of algebraic curves, Γ the underlying metric graph with model $G = (V,E)$, and $\{C_v\}$ the corresponding family of smooth projective curves over κ. Suppose we are given, for each $v \in V$, a non-empty κ-linear subspace F_v of $\kappa(X_v)$. We denote by $\mathcal{F}$ the collection of all F_v. Define the $\mathcal{F}$*-rank* of a divisor $\mathcal{D}$ on $\mathfrak{C}$, denoted $r_{\mathfrak{C},\mathcal{F}}(\mathcal{D})$, to be the maximum integer r such that for any effective divisor $\mathcal{E}$ of degree r, there exists a nonzero rational function $\mathfrak{f} = \{f_\Gamma; \{f_v\}\}$ on $\mathfrak{C}$, with $f_v \in F_v$ for all $v \in V$, such that $\mathcal{D} + \mathrm{div}(\mathfrak{f}) - \mathcal{E} \geq 0$.

The following proposition provides an upper bound for $r_{\mathfrak{C},\mathcal{F}}$ in terms of the maximal dimension of F_v for $v \in V(G)$.

Proposition 6.2. *Let* $s = \max\big(0, \max_{v \in V(G)}(\dim_\kappa(F_v) - 1)\big)$. *Then for any divisor* $\mathcal{D}$ *on* $\mathfrak{C}$, *we have* $r_{\mathfrak{C}X_0,\mathcal{F}}(\mathcal{D}) \leq s$.

Proof. If $\dim_\kappa(F_v) = 0$ for some v, then either $r_{\mathfrak{C},\mathcal{F}}(\mathcal{D}) = 0$ (if $\mathcal{D} \geq 0$) or $r_{\mathfrak{C},\mathcal{F}}(\mathcal{D}) = -1$ (otherwise), and the proposition holds. For the sake of contradiction, suppose now that $r_{\mathfrak{C},\mathcal{F}}(\mathcal{D}) \geq s+1 \geq 1$. Let v be a vertex with $s = \dim_\kappa(F_v) - 1$. For any effective divisor E_v of degree $s+1$ on X_v, there exists a rational function $\mathfrak{f} = \{f_\Gamma; \{f_v\}\}$ on $\mathfrak{C}$ with $f_v \in F_v$ such that $\mathcal{D} - \mathcal{E} + \mathrm{div}(\mathfrak{f}) \geq 0$, where $\mathcal{E} = (s+1)(v) \oplus E_v$. This shows that $D_\Gamma - (s+1)(v) + \mathrm{div}(f_\Gamma) \geq 0$. For a generic choice of E_v, we can assume that the slopes of f_Γ on the edges adjacent to v are fixed [GK] (see also [HMY] and [A, Section 2.1] for a discussion of the cell decomposition of $|D|$ for $D \in \mathrm{Div}(\Gamma)$). For the divisor $D'_v = D_v + \mathrm{div}_v(f_\Gamma)$, this shows that for any generic choice of E_v on X_v, there exists an element $f_v \in F_v$ with $D'_v - E + \mathrm{div}(f_v) \geq 0$. This contradicts the assumption that $\dim_\kappa(F_v) = s+1$, since F_v can define a linear system of projective dimension at most s on X_v. $\square$

6.2.2. *Characterization of limit linear series in terms of the refined rank function.* We retain the terminology from the previous sections. Let X_0 be a curve of compact type, $G = (V, E)$ the dual graph, $\mathfrak{C}X_0$ the regularization of X_0, and let L be a collection of $\mathfrak{g}^r_d$'s $L_v = (\mathcal{L}_v, H_v)$ on X_v, one for each $v \in V$. Fix a root vertex $\mathfrak{r} \in V$ and divisors $\overline{D}_v \in \mathrm{Div}(X_v)$ such that $\mathcal{L}(\overline{D}_v) \sim \mathcal{L}_v$, and let $\mathcal{D}$ be the divisor defined by (7). Denote by $\mathcal{H}$ the family of all $H_v \subseteq H^0(X_v, \mathcal{L}(D_v)) \subset \kappa(X_v)$.

Theorem 6.3. *The following two assertions are equivalent:*

(i) *L is a crude limit $\mathfrak{g}^r_d$ on X_0.*
(ii) *$r_{\mathfrak{C}X_0,\mathcal{H}}(\mathcal{D}) = r$.*

Proof. By the results in Appendix A, in calculating $r_{\mathfrak{C}X_0,\mathcal{H}}(\mathcal{D})$ we can restrict to effective divisors in $\mathrm{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$. In this case, since G is a tree, we can suppose that all the rational functions on Γ which appear in the calculation of $r_{\mathfrak{C}X_0,\mathcal{H}}$ arise via linear interpolation from integral functions $f_G : V \to \mathbb{Z}$. Since G is a tree, the data of an integral function f_G on V is equivalent, up to an additive constant, to an assignment of labels $a_{uv} \in \mathbb{Z}$ to any oriented edge uv of G such that $a_{uv} = -a_{vu}$ for any $\{u, v\} \in E$. (Each edge $\{u, v\}$ gives rise to two oriented edges uv and vu; for any integral function $f_G : V \to \mathbb{Z}$, $a_{uv} = f(u) - f(v)$.)

In what follows, we denote by p_v the unique parent of a vertex $v \neq \mathfrak{r}$ in the rooted tree $G_\mathfrak{r}$. Recall that the κ-point of X_v which corresponds to the edge $\{p_v, v\}$ is denoted by x_v. For any child u of v in $G_\mathfrak{r}$, we denote by y_u the κ-point of X_v which corresponds to the edge $\{v, u\}$.

(i) $\Rightarrow$ (ii): Let L be a crude limit $\mathfrak{g}^r_d$; we want to show that $r_{\mathfrak{C}X_0,\mathcal{H}}(\mathcal{D}) = r$. By Proposition 6.2, we have $r_{\mathfrak{C}X_0,\mathcal{H}}(\mathcal{D}) \leq r$, so it will be enough to prove the reverse inequality $r_{\mathfrak{C}X_0,\mathcal{H}}(\mathcal{D}) \geq r$. Let $\mathcal{E} = E_G \oplus \sum_{v \in V(G)} E_v$ be an effective divisor of degree r in $\mathrm{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$. We need to show the existence of an integral function f_G on V and rational functions $f_v \in H_v \subset \kappa(X_v)$ such that $\mathcal{D} - \mathcal{E} + \mathrm{div}(\mathfrak{f}) \geq 0$, for $\mathfrak{f} = (f_G; f_v, v \in V)$.

The proof of the existence of $f_G, \{f_v\}_{v \in V}$ goes by induction on the distance to $\mathfrak{r}$, starting from the root and going towards the leaves. Roughly speaking, this is done in the following way. If in a given step of the induction we are at some vertex v, then we assign integer values a_{vu} to all the oriented edges vu for any child of u of v in the rooted tree $G_\mathfrak{r}$ and we also choose the rational function f_v on X_v. This is done in a way which will guarantee that at the end of the procedure, for the resulting rational function $\mathfrak{f} = (f_G; f_v, v \in V)$, we have $\mathcal{D} - \mathcal{E} + \mathrm{div}(\mathfrak{f}) \geq 0$.

To describe the details of the edge-labeling procedure, we need to introduce some extra notation. For any vertex v, let $\hat{E}_G(v)$ be the sum of $E_G(v)$ and all the integers $E_G(w)$ for w a descendant of v in the rooted tree $G_\mathfrak{r}$:

$$\hat{E}_G(v) = E_G(v) + \sum_{w \text{ a descendant of } v} E_G(w).$$

We have the recursive equation

$$\hat{E}_G(v) = E_G(v) + \sum_{u \text{ a child of } v} \hat{E}_G(v).$$

We claim that there exists an assignment of integer labels $0 \leq a_{vp_v} \leq d$ to any oriented edge vp_v of G, $v \in V \setminus \mathfrak{r}$, such that Properties (1) and (2) below are satisfied. (Note that $a_{p_v v} = -a_{vp_v}$ and $\{v, p_v\}$ cover all the edges of G.)
For any such assignment, let $\hat{H}_{\mathfrak{r}} = H_{\mathfrak{r}}$ and for any vertex $v \neq \mathfrak{r}$, denote by $\hat{H}_v$ the subspace of H_v which consists of all the global sections of $\mathcal{L}(\overline{D}_v)$ with an order of vanishing at x_v at least equal to $d - a_{uv} \geq 0$.

1. For any vertex v of G, $\hat{H}_v$ has dimension at least $\hat{E}_G(v) + 1$.
2. There exists an element $f_v \in \hat{H}_v$ which has an order of vanishing at least a_{uv} at y_v for any child u of v.

Note that by the first property (1) above, if $v \neq \mathfrak{r}$, f_v has an order of vanishing at least $d - a_{vp_v}$ at x_v. Recall that $D_v = \overline{D}_v - d(x_v)$ and $a_{vu} = -a_{uv}$. For $v = \mathfrak{r}$, there is no $p_{\mathfrak{r}}$ and we let $a_{vp_v} = 0$. Therefore, for any vertex v, the two properties (1) and (2) above imply that

$$D_v + a_{vp_v}(x_v) + \sum_{u \text{ a child of } v} a_{vu}\,(y_u) + \operatorname{div}(f_v) \geq 0.$$

It follows that for the resulting rational function $\mathfrak{f}$ on $\mathfrak{C}X_0$, we will have $\mathcal{D} - \mathcal{E} + \operatorname{div}(\mathfrak{f}) \geq 0$ as desired. So we are left to prove the claim.

We provide a constructive proof by inductively assigning values to a_{vp_v}, starting from the root and going down toward the leaves. Let v be a vertex of G such that a_{vp_v} has been assigned a value and $\hat{H}_v$ has dimension at least $\hat{E}_G(v) + 1$. Note that for $v = \mathfrak{r}$, since $\hat{E}_G(\mathfrak{r}) = \deg(\mathcal{E}) = r$ and $\dim_\kappa(H_{\mathfrak{r}}) = r+1$, we have $\dim_\kappa(\hat{H}_{\mathfrak{r}}) = \hat{E}_G(\mathfrak{r}) + 1$. Thus the base of the induction holds trivially. We now show the existence of an assignment of non-negative integers $0 \leq a_{uv} \leq d$ to the oriented edges uv for any child u of v, and the existence of an element $f_v \in \hat{H}_v \subset H^0(X_v, \mathcal{L}(D_v))$ such that:

- Property (1) holds for any child u of v, i.e., $\hat{H}_u$ has dimension at least $\hat{E}_G(u) + 1$, where $\hat{H}_u$ is the subspace of H_u consisting of all global sections of $\mathcal{L}(D_u)$ whose order of vanishing at x_u is at least $d - a_{uv}$.
- Property (2) holds for v and f_v, i.e., f_v has order of vanishing at least a_{uv} at y_u for any child u of v.

To construct such an assignment of values $0 \leq a_{uv} \leq d$, we proceed as follows. For each child u of v, by Condition (6) applied to u and v, we have

$$a_i^{L_v}(y_u) + a_{r-i}^{L_u}(x_u) \;\geq\; d$$

for all $0 \leq i \leq r$. In particular, for $i = \hat{E}_G(u)$, we have

$$a^{L_u}_{r - \hat{E}_G(u)}(x_u) \geq d - a^{L_v}_{\hat{E}_G(u)}(y_u).$$

Define $a_{uv} := a^{L_v}_{\hat{E}_G(u)}(y_u)$, so that the above inequality reads $a^{L_u}_{r-\hat{E}_G(u)}(x_u) \geq d - a_{uv}$. By the definition of the sequence $a_i^{L_u}$, this shows that the space of sections in H_u whose order of vanishing at x_u is at least $d - a_{uv}$ has dimension at least $\hat{E}_G(u) + 1$; thus Property (1) is satisfied for any child u of v. It remains to show the second condition, namely the existence of $f_v \in \hat{H}_v$ with the required orders of vanishing at the points y_u. By the definition of the sequence $a_i^{L_v}$, the subspace of sections in H_v whose order of vanishing at y_u is at least $a_{uv} = a^{L_v}_{\hat{E}_G(u)}(y_u)$ has codimension at most $\hat{E}_G(u)$ in H_v. *A posteriori*, the subspace of sections in $\hat{H}_v$ with an order of vanishing at least a_{uv} at y_u has codimension at most $\hat{E}_G(u)$. Therefore the space of sections in $\hat{H}_v$ with the required orders of vanishing at each point y_u for children u of v has codimension at most $\sum_{u \text{ a child of } v} \hat{E}_G(u) \leq \hat{E}_G(v)$. Since $\dim_\kappa \hat{H}_v \geq \hat{E}_G(v) + 1$, we infer the existence of f_v satisfying Property (2).

(ii) $\Rightarrow$ (i): Suppose that $r_{\mathfrak{C},\mathcal{H}}(\mathcal{D}) = r$. We first observe that for any other choice of a vertex $\mathfrak{r}'$ as root, if $\mathcal{D}'$ denotes the corresponding divisor on $\mathfrak{C}$, then $\mathcal{D}$ and $\mathcal{D}'$ differ by a principal divisor $\mathrm{div}(\mathfrak{f})$ for a rational function $\mathfrak{f} = (f_G; f_v = 1,\, v \in V)$ on $\mathfrak{C}X_0$. Thus $r_{\mathfrak{C},\mathcal{H}}(\mathcal{D}') = r$ as well. This shows that the choice of $\mathfrak{r}$ is arbitrary, and so it will be enough to prove the validity of Condition (6) for $v = \mathfrak{r}$ and a child u of $\mathfrak{r}$. Let $0 \leq i \leq d$. For any effective divisor $\mathcal{E} = ((r-i)(v) + i(u)) \oplus (E_v + E_u)$, where E_v and E_u are two effective divisors of degree $r-i$ and i on X_v and X_u, respectively, there must exist a rational function $\mathfrak{f} = (f_G;\, f_w, w \in V))$, $f_w \in H_w$, such that $\mathcal{D} - \mathcal{E} + \mathrm{div}(\mathfrak{f}) \geq 0$. For any two adjacent v ertices w and z in G, let $a_{wz} = f(w) - f(z)$ as before. Since for any vertex $w \neq \mathfrak{r}$, D_w has degree zero, a simple induction starting from the leaves and going towards the root shows that $a_{wp_w} \geq 0$. This in particular implies that $D_\mathfrak{r} + \mathrm{div}(f_\mathfrak{r}) - E_\mathfrak{r} - a_{u\mathfrak{r}}(y_u) \geq 0$, and $D_u + \mathrm{div}(f_u) - E_u - (d - a_{u\mathfrak{r}})(x_u) \geq 0$. For a generic choice of $E_\mathfrak{r}$ and E_u, we can assume that $a_{u\mathfrak{r}} = a$ is a constant. In other words, there exists an integer $0 \leq a \leq d$ such that the sublinear system of $H_\mathfrak{r}$ (resp., H_u) which consists of those sections with an order of vanishing at least a (resp. $d-a$) at y_u (resp. x_u) has projective dimension at least $r-i$ (resp. i). By the definition of the sequences $a^{L_\mathfrak{r}}(.)$ and $a^{L_u}(.)$, this simply means that $a_i^{L_\mathfrak{r}}(y_u) \geq a$ and $a_{r-i}^{L_u}(x_u) \geq d-a$, and thus, $a_i^{L_\mathfrak{r}}(y_u) + a_{r-i}^{L_u}(x_u) \geq d$, which is Condition (6) for the two irreducible components $X_\mathfrak{r}$ and X_u of X_0. ☐

Corollary 6.4. *Let κ be of characteristic zero. Then L is a refined limit series iff $r_{\mathfrak{C},\mathcal{H}}(\overline{D}) = r$ and all the ramification points of L_v are smooth points of X_v for all $v \in V(G)$.*

Proof. This is a direct consequence of Theorem 6.3 above and the Plücker formula, c.f. [EH, Proposition 1.1]. ☐

We note that the *a priori* dependence of the family $\mathcal{H}$ on the X_v-aspects L_v can be removed, and recovered from the condition on the rank, in the following sense. For any crude limit $\mathfrak{g}^r_d$ L on X_0 with X_v-aspect $L_v = (\mathcal{L}_v, H_v)$, choose a divisor $D_v \in \mathrm{Div}^d(X_v)$ with $\mathcal{L}_v \simeq \mathcal{L}(D_v)$. Two divisors $\mathcal{D}$ and $\mathcal{D}'$ in $\mathrm{Div}(\mathfrak{C}X_0)_\mathbb{Z}$ are called *combinatorially equivalent* if they differ by the divisor of a rational function on $\mathfrak{C}X_0$ of the form $\mathfrak{f} = (f_\Gamma, (f_v)_{v \in V})$ with all f_v constant.

The following result is essentially a reformulation of Theorem 1.9 from the Introduction and easily implies the result stated there.

Theorem 6.5. *Let X_0 be a curve of compact type and $\mathfrak{C}X_0$ the regularization of X_0. Then there is a bijective correspondence between the following:*

- *Pairs $(L, \{D_v\})$ consisting of a crude limit $\mathfrak{g}^r_d$ L on X_0 and a collection of divisors D_v on X_v with $L_v \simeq \mathcal{L}(D_v)$ for any irreducible component X_v of X_0; and*
- *Pairs $(\mathcal{H}, [\mathcal{D}])$, where $\mathcal{H} = \{H_v\}$, H_v is an $(r+1)$-dimensional subspace of $\kappa(X_v)$ for each $v \in V$, and $[\mathcal{D}]$ is the combinatorial linear equivalence class of a divisor $\mathcal{D} \in \mathrm{Div}(\mathfrak{C}X_0)_\mathbb{Z}$ of degree d on $\mathfrak{C}X_0$ with $r_{\mathfrak{C}X_0,\mathcal{H}}(\mathcal{D}) = r$.*

In particular, the data of a crude limit $\mathfrak{g}^r_d$ on X_0 is equivalent to the data of a pair $(\mathcal{H}, [\mathcal{D}])$ with $r_{\mathfrak{C}X_0,\mathcal{H}}(\mathcal{D}) = r$.

Proof. We already know by Theorem 6.3 that any crude limit $\mathfrak{g}^r_d$ gives a pair $(\mathcal{H}, [\mathcal{D}])$ with $r_{\mathfrak{C}X_0,\mathcal{H}}(\mathcal{D}) = r$.

Let $\mathcal{H}$ be a collection of $(r+1)$-dimensional subspaces $H_v \subset \kappa(X_v)$ for $v \in V$, and let $\mathcal{D}$ be a divisor of degree d in $\mathrm{Div}(\mathfrak{C}X_0)_{\mathbb{Z}}$ with $r_{\mathfrak{C}X_0,\mathcal{H}}(\mathcal{D}) = r$. We show the existence of a $\mathfrak{g}^r_d$ $\mathcal{L}_v = (L_v, H_v)$ on X_v for each $v \in V$ such that the collection $L = \{L_v\}$ defines a crude limit $\mathfrak{g}^r_d$ on X_0, and such that $\mathcal{D}$ is (combinatorially) linearly equivalent to the divisor $\mathcal{D}^L$ associated to L by (7).

Write $\mathcal{D} = D_\Gamma \oplus \sum_{v\in V(G)} D_v$. Since Γ has genus zero, for any vertex v in $V(G)$ there exists a rational function $f_{\Gamma,v}$ on Γ such that $D_\Gamma + \mathrm{div}(f_{\Gamma,v}) = d(v)$. Define $\widetilde{D}_v := D_v + \mathrm{div}_v(f_{\Gamma,v})$. Let $\mathcal{L}_v := \mathcal{L}(\widetilde{D}_v)$. We claim that $H_v \subset H^0(X_v, \mathcal{L}_v)$; this amounts to showing that $\mathrm{div}(f_v) + \widetilde{D}_v \geq 0$ for any $f_v \in H_v$. If we denote by H'_v the subspace of H_v consisting of all rational functions in H_v with this property, it will be enough show that $\dim_\kappa(H'_v) = r+1$.

First, we note that for any rational function f_Γ on Γ with $D_\Gamma + \mathrm{div}(f_\Gamma) \geq 0$, we have $\mathrm{div}_v(f_{\Gamma,v}) - \mathrm{div}_v(f_\Gamma) \geq 0$. Indeed, $D_\Gamma + \mathrm{div}(f_\Gamma) \geq 0$ implies that for any edge e adjacent to v, the slope of f_Γ in the tangent direction corresponding to e is bounded above by the sum of the coefficients $D_\Gamma(w)$ of the points w lying on the connected component of $\Gamma \setminus \{v\}$ which contains the edge e. This sum is precisely the slope of $f_{\Gamma,v}$ in the tangent direction corresponding to e.

For any effective divisor E_v of degree r on X_v, let $\mathcal{E}_v = r(v) \oplus E_v$ be a corresponding divisor on $\mathfrak{C}X_0$. Since $r_{\mathfrak{C}X_0,\mathcal{H}}(\mathcal{D}) = r$, there exists a rational function $\mathfrak{f} = \{f_\Gamma; f_w, w \in V(G)\}$ such that $\mathrm{div}(\mathfrak{f}) + \mathcal{D} - \mathcal{E}_v \geq 0$. Since $\mathrm{div}(f) + D \geq 0$, we have $\widetilde{D}_v = D_v + \mathrm{div}_v(f_{\Gamma,v}) \geq D_v + \mathrm{div}_v(f_\Gamma)$. By the definition of $\mathrm{div}(\mathfrak{f})$ on $\mathfrak{C}X_0$, this implies that $\mathrm{div}(f_v) + \widetilde{D}_v - E_v \geq 0$, in particular f_v belongs to H'_v. Since this holds for any E_v, we infer that the linear series on X_v defined by H'_v has (projective) dimension r and so $H'_v = H_v$. We conclude that $L_v = (\mathcal{L}_v, H_v)$ is a $\mathfrak{g}^r_d$ on X_v, and thus the collection of L_v for $v \in V$ defines a crude limit L on X_0.

Let $\mathcal{D}^L$ be the divisor associated to L by (7). By the explicit definition of $\mathcal{D}^L$ in terms of $\widetilde{D}_v$ and a root $\mathfrak{r} \in V$, it is straightforward to see that $\mathcal{D}^L$ and $\mathcal{D}$ differ by the divisor of a rational function $\mathfrak{f} = (f_\Gamma; f_v = 1, v \in V(G))$ on $\mathfrak{C}X_0$ for a rational function f_Γ on Γ defined by an integral function $f : V \to \mathbb{Z}$. $\square$

6.3. **Limit linear series for metrized complexes.** Let $\mathfrak{C}$ be a metrized complex of algebraic curves over κ, Γ the underlying metric graph with model $G = (V, E)$, and $\{C_v\}$ the corresponding collection of smooth projective curves over κ.

We define a *(crude) limit* $\mathfrak{g}^r_d$ on $\mathfrak{C}$ to be the data of a pair $(\mathcal{H}, \mathcal{D})$ consisting of a divisor $\mathcal{D}$ of degree d on $\mathfrak{C}$ and a collection $\mathcal{H}$ of $(r+1)$-dimensional subspaces $H_v \subset \kappa(C_v)$, for $v \in V(G)$, such that $r_{\mathfrak{C},\mathcal{H}}(\mathcal{D}) = r$.

Remark 6.6. Since $r_{\mathfrak{C},\mathcal{H}}(\mathcal{D}) \leq r_{\mathfrak{C}}(\mathcal{D})$ for all $\mathcal{H}$ as above and all $\mathcal{D} \in \mathrm{Div}(\mathfrak{C})$, it follows immediately from Theorem 3.2 that for any limit $\mathfrak{g}^r_d$ on $\mathfrak{C}$ we have if $d \geq 2g$ or $r \geq g$, then $r + g \leq d$ (Riemann's inequality), and if $\mathcal{D}$ is *special* then $r \leq d/2$ (Clifford's inequality). For curves of compact type, these inequalities are established (by a completely different proof) in [EH].

Theorem 6.7. *Let X be a smooth proper curve over $\mathbb{K}$, $\mathfrak{X}$ a semistable model for X, and let $\mathfrak{C}\mathfrak{X}$ be the metrized complex associated to $\mathfrak{X}$. Let D a divisor on X and let $L_\eta = (\mathcal{L}(D), H_\eta)$,* *for $H_\eta \subseteq H^0(X,\mathcal{L}(D)) \subset \mathbb{K}(X)$, be a $\mathfrak{g}^r_d$ on X. For any vertex $v \in V$, define H_v as the κ-vector space defined by the reduction to $\kappa(C_v)$ of all the rational functions in H_η (c.f. Section 4.11), and let $\mathcal{H} = \{H_v\}_{v\in V}$. Then the pair $(\tau_*^{\mathfrak{CX}}(D), \mathcal{H})$ is a limit $\mathfrak{g}^r_d$ on $\mathfrak{CX}$.*

Proof. By Lemma 4.3, all the subspaces H_v have dimension $r+1$ over κ. By Proposition 6.2, we only have to show that $r_{\mathfrak{CX},\mathcal{H}} \geq r$. By the results in Appendix A, it suffices to show that for any effective divisor $\mathcal{E} \in \mathrm{Div}(\mathfrak{CX})$ of degree r supported on a subset $R = \bigcup_v R_v$ of $\mathfrak{CX}$, with $R_v \subset C_v(\kappa) \setminus \mathcal{A}_v$ of size $g_v + 1$, there exists a rational function $\mathfrak{f} = (f_\Gamma; f_v, v \in V(G))$ on $\mathfrak{CX}$ such that $f_v \in H_v$ for all $v \in V$ and $\mathrm{div}(\mathfrak{f}) + \tau_*^{\mathfrak{CX}}(D) - \mathcal{E} \geq 0$. For any such $\mathcal{E}$, there exists an effective divisor of degree r on X such that $\tau_*^{\mathfrak{CX}}(E) = \mathcal{E}$. Since L_η is a $\mathfrak{g}^r_d$ on X, there exists a rational function $f \in H_\eta$ such that $D - E + \mathrm{div}(f) \geq 0$. Let $\mathfrak{f} = (f_\Gamma; \widetilde{f_v}, v \in V(G))$ be the rational function on $\mathfrak{CX}$ defined by f. We conclude that $\tau_*^{\mathfrak{CX}}(D - E + \mathrm{div}(f)) = \tau_*^{\mathfrak{CX}}(D) - \mathcal{E} + \tau_*^{\mathfrak{CX}}(\mathrm{div}(f)) = \tau_*^{\mathfrak{CX}}(D) - \mathcal{E} + \mathrm{div}(\mathfrak{f}) \geq 0$. Since $\widetilde{f_v} \in H_v$, this shows that $r_{\mathfrak{CX},\mathcal{H}} \geq r$ as desired. ☐

For the regularization $\mathfrak{C}X_0$ of a strongly semistable curve X_0 over κ, Theorem 6.7 can be reformulated as follows. Let $\phi : \mathcal{X} \to \operatorname{Spec} R$ be a regular smoothing of X_0 over a discrete valuation ring R and let $L_\eta = (\mathcal{L}_\eta, H_\eta)$ be a $\mathfrak{g}^r_d$ over the generic fiber $\mathcal{X}_\eta$ of ϕ. Consider an extension $\mathcal{L}$ of L_η to $\mathcal{X}$, and let $\mathcal{L}_0$ be the restriction of $\mathcal{L}$ to X_0. Let $\mathcal{D}$ be a divisor on $\mathfrak{C}X_0$ in the linear equivalence class of $\pi(\mathcal{L}_0)$, where $\pi : \mathrm{Pic}(X_0) \to \oplus_{v\in V(G)} \mathrm{Pic}(X_v)$ is as in Section 2. There exists a family $\mathcal{H} = \{H_v\}_{v \in V(G)}$ of $(r+1)$-dimensional subspaces $H_v \subset \kappa(X_v)$ such that $r_{\mathfrak{C}X_0,\mathcal{H}}(\mathcal{D}) = r$. In addition, H_v can be explicitly constructed as follows. For any vertex v, let $f_{G,v} : V(G) \to \mathbb{Z}$ be the integral function such that for the rational function $\mathfrak{f}_v = (f_{G,v}; f_u = 1,\, u \in V(G))$ on $\mathfrak{C}X_0$, the divisor $\mathcal{D} + \mathrm{div}(f_{G,v})$ is v-reduced. Consider the extension $\mathcal{L}\left(\sum_v f_{G,v}(u) X_u\right)$ of $\mathcal{L}_\eta$ to $\mathcal{X}$ and let $H_{f_{G,v}}$ be the closure of H_η in $H^0\big(\mathcal{X}, \mathcal{L}\left(\sum_v f_{G,v}(u) X_u\right)\big)$. Define $H_v := H_{f_{G,v}}|_{X_v}$ and $\tau_*(L_\eta) := (\mathcal{D}, \mathcal{H})$. Finally, for any $v \in V$ let $\mathcal{D}^v = D^v_\Gamma \oplus \sum_{u\in V} D^v_u$ be the v-reduced divisor linearly equivalent to $\mathcal{D}$, and define $d_v := \mathcal{D}_\Gamma(v) = \deg(D^v_v)$.

Theorem 6.8. *Let X_0 be a strongly semistable curve over κ and $\mathfrak{C}X_0$ the regularization of X_0. For any regular smoothing $\pi : X \to \operatorname{Spec} R$ of X_0 over a discrete valuation ring R, $\tau_*(L_\eta)$ as defined above is a limit $\mathfrak{g}^r_d$ on X_0. In addition, for each vertex $v \in V(G)$ we have $H_v \subseteq H^0(X_v, \mathcal{L}(D^v_v))$, so that $(\mathcal{L}(D^v_v), H_v)$ defines a $\mathfrak{g}^r_{d_v}$ on the irreducible component X_v of X_0.*

We conclude this section with the following strengthening of Theorem 6.7 which will be used in §6.5.

Theorem 6.9. *Let X be a smooth proper curve over $\mathbb{K}$, $\mathfrak{X}$ a semistable model for X, and let $\mathfrak{CX}$ be the metrized complex associated to $\mathfrak{X}$. Let D be a divisor on X and set $\mathcal{D} = \tau_*^{\mathfrak{CX}}(D)$. Let $L_\eta = (\mathcal{L}(D), H_\eta)$ be a $\mathfrak{g}^r_d$ on X and let $\mathcal{H} = \{H_v\}_{v\in V}$ be as in the statement of Theorem 6.7. Let E° be an effective divisor of degree e supported on the smooth locus of the special fiber of $\mathfrak{X}$, and define $\mathcal{E}^\circ = (E^\circ_\Gamma, \{E^\circ_v\}) \in \mathrm{Div}(\mathfrak{CX})$ by letting E°_v be the restriction of E° to C_v and setting $E^\circ_\Gamma = \sum_{v\in V} \deg(E^\circ_v)(v)$. Suppose that $H_v \subseteq L(D_v - E^\circ_v)$ for all v. Then the pair $(\mathcal{D} - \mathcal{E}^\circ, \mathcal{H})$ is a limit $\mathfrak{g}^r_{d-e}$ on $\mathfrak{CX}$.*

Proof. In the view of the results in Appendix A, the proof is similar to the proof of Theorem 6.7, by requiring each subset $R_v \subset C_v(\kappa) \setminus \mathcal{A}_v$ to be disjoint from the support of E°_v. For any effective divisor $\mathcal{E}$ of degree r supported on $R = \bigcup R_v$, the proof of Theorem 6.7 shows that there exists a rational function $\mathfrak{f} = (f_\Gamma; \widetilde{f}_v, v \in V(G))$ such that $\mathcal{D} - \mathcal{E} + \mathrm{div}(\mathfrak{f}) \geq 0$. We claim that in fact $\mathcal{D} - \mathcal{E} + \mathrm{div}(\mathfrak{f}) - \mathcal{E}^\circ \geq 0$, which implies that $(\mathcal{D} - \mathcal{E}^\circ, \mathcal{H})$ is a limit $\mathfrak{g}^r_{d-e}$ on $\mathfrak{CX}$ as desired. To see this, note that since $E^\circ_\Gamma = \sum_{v \in V} \deg(E^\circ_v)(v)$ and $\mathrm{div}(f_\Gamma) \geq E_\Gamma - D_\Gamma$ by hypothesis, it suffices to prove that the v-component of $\mathrm{div}(\mathfrak{f})$ is at least the v-component of $\mathcal{E} + \mathcal{E}^\circ - \mathcal{D}$ for all $v \in V$, i.e.,

$$\mathrm{div}(\mathfrak{f})_v := \mathrm{div}(\widetilde{f}_v) + \mathrm{div}_v(f_\Gamma) \geq E_v + E^\circ_v - D_v. \tag{8}$$

If $z \in \mathrm{supp}(E^\circ_v)$, then since $\mathrm{supp}(E^\circ_v)$ is disjoint from $\mathcal{A}_v \supset \mathrm{supp}(\mathrm{div}_v(f_\Gamma))$ we have $\mathrm{div}(\mathfrak{f})_v(z) = \mathrm{div}(\widetilde{f}_v)(z)$, and since $\mathrm{supp}(E^\circ_v)$ is disjoint from $R_v \supset \mathrm{supp}(E_v)$ we have $E_v(z) = 0$. Thus (8) holds in this case since $H_v \subset L(D_v - E^\circ_v)$ implies that $\mathrm{div}(\widetilde{f}_v)(z) \geq E^\circ_v(z) - D_v(z)$. On the other hand, if $z \notin \mathrm{supp}(E^\circ_v)$, then $E^\circ_v(z) = 0$ and thus (8) holds because by assumption we have $\mathrm{div}(\mathfrak{f}) \geq \mathcal{E} - \mathcal{D}$ and in particular $\mathrm{div}(\mathfrak{f})_v(z) \geq E_v(z) - D_v(z)$.

Thus (8) is valid for all $z \in C_v(\kappa)$ as desired. □

6.4. **Remarks on completion and smoothing.** Given a divisor $\mathcal{D}$ of degree d and rank r on $\mathfrak{C}$, one may ask if there exists a collection $\mathcal{H} = \{H_i\}$, with H_i an $(r+1)$-dimensional subspace of $\kappa(C_i)$ for all i, such that $r_{\mathfrak{C},\mathcal{H}}(\mathcal{D}) = r$. (Recall that we always have $r_{\mathfrak{C},\mathcal{H}}(\mathcal{D}) \leq r_{\mathfrak{C}}(\mathcal{D})$.) If so, we say that $\mathcal{D}$ can be *completed to a limit* $\mathfrak{g}^r_d$ on $\mathfrak{C}$. Moreover, Theorem 6.7 shows that if X is a smooth proper curve over $\mathbb{K}$, (D, L_η) is a $\mathfrak{g}^r_d$ on X, and $\mathfrak{CX}$ is the metrized complex associated to a semistable model $\mathfrak{X}$ for X, then the associated pair $(\tau^{\mathfrak{CX}}_*(D), \mathcal{H})$ (which we call the *specialization* of $(\mathcal{D}, \mathcal{H})$) is a limit $\mathfrak{g}^r_d$ on $\mathfrak{CX}$. We say that a limit $\mathfrak{g}^r_d$ $(\mathcal{D}, \mathcal{H})$ on a metrized complex $\mathfrak{C}$ over κ is *smoothable* if there exists a smooth proper curve X over $\mathbb{K}$ such that $\mathfrak{C} = \mathfrak{CX}$ for some semistable model $\mathfrak{X}$ of X and $(\mathcal{D}, \mathcal{H})$ arises the specialization of a limit $\mathfrak{g}^r_d$ on X. Note that by the results in [ABBR], every $\mathrm{val}(\mathbb{K}^*)$-rational metrized complex $\mathfrak{C}$ over κ is of the form $\mathfrak{C} = \mathfrak{CX}$ for some smooth proper curve X over $\mathbb{K}$ and some semistable model $\mathfrak{X}$ of X, where $\mathrm{val}(\mathbb{K}^*)$-*rational* means that the edge lengths in G all belong to the value group of $\mathbb{K}$.

In this section, we show by example that not every divisor $\mathcal{D}$ of degree d and rank r on a metrized complex $\mathfrak{C}$ can be completed to a limit $\mathfrak{g}^r_d$, and not every limit $\mathfrak{g}^r_d$ on a $\mathrm{val}(\mathbb{K}^*)$-rational metrized complex $\mathfrak{C}$ is smoothable.

Example 6.10 (A degree d rank r divisor which cannot be completed to a limit $\mathfrak{g}^r_d$). Consider the weighted graph Γ depicted in Figure 1 (with arbitrary edge lengths), and let $\mathfrak{C}$ be a metrized complex whose underlying weighted metric graph is Γ. For any point p in $C_v \simeq \mathbb{P}^1$, consider the degree 2 divisor $\mathcal{D} = 2(v) \oplus 2(p)$. We claim that $\mathcal{D}$ has rank 1 on $\mathfrak{C}$ but cannot be completed to a limit $\mathfrak{g}^1_2$. The fact that $\mathcal{D}$ has rank 1 is straightforward and left to the reader. Assume for the sake of contradiction that there exists a limit $\mathfrak{g}^1_2$ $(\mathcal{H}, \mathcal{D})$ on $\mathfrak{C}$. Let x_1, x_2, x_3 (resp. y_1, y_2, y_3) be the points of $C_v = \mathbb{P}^1$ (resp. $C_i := C_{v_i}$) corresponding to the three edges of Γ. Fix $i \in \{1, 2, 3\}$, choose a point $q_i \neq y_i$ in C_i and let $\mathcal{E} = (v_i) \oplus (q_i)$. For any rational function $\mathfrak{f}$ on $\mathfrak{C}$ such that $\mathcal{D} - \mathcal{E} + \mathrm{div}(\mathfrak{f}) \geq 0$, the fact that C_i has genus one implies that $\mathrm{div}_{v_i}(f_\Gamma) \geq 2(y_i)$, since there is no rational function on C_i whose divisor is $(q_i) - (y_i)$. Therefore $\mathrm{div}(f_\Gamma) = 2(v_i) - 2(v)$. Restricting $\mathfrak{f}$ to C_v gives a rational function f_i on C_v with $\mathrm{div}(f_i) = 2(x_i) - 2(p)$. Since $(\mathcal{H}, \mathcal{D})$ is by assumption a limit $\mathfrak{g}^1_2$, we must have $f_1, f_2, f_3 \in H_v$. But a simple calculation shows that the f_i are linearly independent, implying that $\dim_\kappa(H_v) \geq 3$, a contradiction.

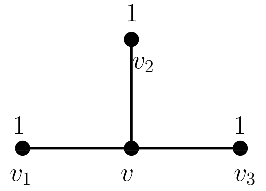

FIGURE 1. A weighted metric graph underlying a metrized complex possessing a degree 2 rank 1 divisor which cannot be completed to a limit $\mathfrak{g}^1_2$.

Remark 6.11. If the residue field κ of $\mathbb{K}$ has characteristic different from 2, any hyperelliptic metric graph which cannot be realized as the skeleton of a hyperelliptic curve over $\mathbb{K}$ provides an example of a divisor of degree two and rank one on a metrized complex $\mathfrak{C}$ over k (with all curves C_v isomorphic to $\mathbb{P}^1$) which cannot be completed to a limit $\mathfrak{g}^1_2$. (There are many examples of this phenomenon presented in [ABBR], which also contains a complete characterization of all hyperelliptic metric graphs which can be lifted to a hyperelliptic curve over $\mathbb{K}$.) Indeed, it is possible to show that any limit $\mathfrak{g}^1_2$ on a metrized complex $\mathfrak{C}$ provides a finite harmonic morphism of degree two from $\mathfrak{C}$ to a metrized complex of genus zero (in the terminology of [ABBR]). By the main lifting result in [ABBR], such a finite harmonic morphism can be lifted to a degree two map from a smooth proper curve $X/\mathbb{K}$ with associated metrized complex $\mathfrak{C}$ to $\mathbb{P}^1$, which implies that the underlying metric graph of $\mathfrak{C}$ is the skeleton of a hyperelliptic curve over $\mathbb{K}$.

Example 6.12 (A limit $\mathfrak{g}^r_d$ which is not smoothable). This phenomenon occurs already in the compact type case: the original paper of Eisenbud-Harris on limit linear series [EH] provides an example (Example 3.2 of *loc. cit.*) of a limit $\mathfrak{g}^2_4$ which cannot be smoothed; we recall their example here.

The metrized complex $\mathfrak{C}$ in question is associated to a semistable curve X_0 with two components meeting transversely at a point p; one of the components is a hyperelliptic curve Y of genus at least 4 and the other component Z has genus zero. The point p is a ramification point of the $\mathfrak{g}^1_2$ on Y, and without loss of generality we may assume that p corresponds to the point 0 on $\mathbb{P}^1 \cong Z$. The dual graph G of X_0 has two vertices v_Y and v_Z connected by a single edge, which may take to be of length one in Γ. Now define a divisor $\mathcal{D}$ on $\mathfrak{C}$ by setting $D_\Gamma = 4(v_Z), D_Z = 4(\infty)$, and $D_Y = 0$. If we define $\mathcal{H} = \{H_Y, H_Z\}$, where $H_Y = H^0(Y, 4(p))$ and H_Z is the span of $1, t^2 + t^3, t^4$ (with t a local parameter at 0 on $\mathbb{P}^1$), then by [EH, Example 3.2] and Theorem 6.3 above, the pair $(\mathcal{D}, \mathcal{H})$ defines a limit $\mathfrak{g}^2_4$ on $\mathfrak{C}$. (It is also not hard to show this directly using our definitions.)

Eisenbud and Harris give the following argument showing that this limit $\mathfrak{g}^2_4$ cannot be smoothed. The incomplete linear series $(4(\infty), H_Z)$ on Z is base-point free and defines a morphism $\phi : Z \to \mathbb{P}^2$ which is birational onto its image. If $(\mathcal{D}, \mathcal{H})$ could be smoothed then we could view it as the generic fiber of a degree 4 morphism $\phi_\eta : X \to \mathbb{P}^2$ which is also birational onto its image, but then $\phi_\eta(X)$ would be a plane quartic of geometric genus at least 4, which is absurd.

Remark 6.13. Despite examples like Example 6.12 above, Eisenbud and Harris also prove some positive results about smoothing, e.g. Theorem 3.4 in *loc. cit.*. It would be very interesting to prove similar results for limit linear series on curves of non-compact type in the sense of the present paper.

6.5. **Application to bounding the number of rational points on curves.** In this section, we explain how limit linear series on metrized complexes of curves can be used to illuminate the proof of a recent theorem due to E. Katz and D. Zureick-Brown [KZB] and put it into a broader context.

Let K be a number field and suppose X is a smooth, proper, geometrically integral curve over K of genus $g \geq 2$. Let J be the Jacobian of X, which is an abelian variety of dimension g defined over K. If the Mordell-Weil rank r of $J(K)$ is less than g, Coleman [Co] adapted an old method of Chabauty to prove that if $p > 2g$ is a prime which is unramified in K and $\mathfrak{p}$ is a prime of good reduction for X lying over p, then $\#X(K) \leq \#\bar{X}(\mathbf{F}_\mathfrak{p}) + 2g - 2$. Here $\bar{X}$ denotes the special fiber of a smooth proper model for X over the completion $\mathcal{O}_\mathfrak{p}$ of $\mathcal{O}_K$ at $\mathfrak{p}$ and $\mathbf{F}_\mathfrak{p} = \mathcal{O}_K/\mathfrak{p}$. Stoll [St] improved this bound by replacing $2g-2$ with $2r$. Lorenzini and Tucker [LT] (see also [MP]) proved the same bound as Coleman without assuming that X has good reduction at $\mathfrak{p}$; in their bound, $\bar{X}(\mathbf{F}_\mathfrak{p})$ is replaced by $\bar{\mathfrak{X}}^{\mathrm{sm}}(\mathbf{F}_\mathfrak{p})$ where $\mathfrak{X}$ is a proper regular model for X over $\mathcal{O}_\mathfrak{p}$ and $\bar{\mathfrak{X}}^{\mathrm{sm}}$ is the smooth locus of the special fiber of $\mathfrak{X}$. Katz and Zureick-Brown combine the improvements of Stoll and Lorenzini-Tucker by proving:

Theorem 6.14. *Let K be a number field and suppose X is a smooth, proper, geometrically integral curve over K of genus $g \geq 2$. Suppose the Mordell-Weil rank r of $J(K)$ is less than g, and that $p > 2g$ is a prime which is unramified in K. Let $\mathfrak{p}$ be a prime of $\mathcal{O}_K$ lying over p and let $\mathfrak{X}$ be a proper regular model for X over $\mathcal{O}_\mathfrak{p}$. Then*

$$\#X(K) \leq \#\bar{\mathfrak{X}}^{\mathrm{sm}}(\mathbf{F}_\mathfrak{p}) + 2r.$$

In order to explain the main new idea in the paper of Katz and Zureick-Brown, we first quickly recall the basic arguments used by Coleman, Stoll, and Lorenzini-Tucker. (See [MP] for a highly readable and more detailed overview.) Assume first that we are in the setting of Coleman's paper, so that $r < g$, $p > 2g$ is a prime which is unramified in K, and X has good reduction at the prime $\mathfrak{p}$ lying over p. Fix a rational point $P \in X(K)$ (if there is no such point, we are already done!). Coleman associates to each regular differential ω on X over $K_\mathfrak{p}$ (the $\mathfrak{p}$-adic completion of K) a "definite p-adic integral" $\int_P^Q \omega \in K_\mathfrak{p}$. If V_{chab} denotes the vector space of all ω such that $\int_P^Q \omega = 0$ for all $Q \in X(K)$, Coleman shows that $\dim V_{\mathrm{chab}} \geq g - r > 0$. Locally, p-adic integrals are obtained by formally integrating a power series expansion for ω with respect to a local parameter. Using this observation and an elementary Newton polygon argument, Coleman proves that

$$\#X(K) \leq \sum_{\widetilde{Q} \in \bar{X}(\mathbf{F}_\mathfrak{p})} \left(1 + n_{\widetilde{Q}}\right),$$

where $n_{\widetilde{Q}}$ is the minimum over all nonzero ω in V_{chab} of $\mathrm{ord}_{\widetilde{Q}} \widetilde{\omega}$; here $\widetilde{\omega}$ denotes the reduction of ω to $\bar{X}$. If we choose any nonzero $\omega \in V_{\mathrm{chab}}$, then the fact that the canonical divisor class on $\bar{X}$ has degree $2g-2$ gives

$$\sum_{\widetilde{Q} \in \bar{X}(\mathbf{F}_\mathfrak{p})} n_{\widetilde{Q}} \leq \sum_{\widetilde{Q} \in \bar{X}(\mathbf{F}_\mathfrak{p})} \mathrm{ord}_{\widetilde{Q}} \widetilde{\omega} \leq 2g-2,$$

which yields Coleman's bound.

Stoll observed that one could do better than this by adapting the differential ω to the point $\widetilde{Q}$ rather than using the same differential ω on all residue classes. Define the *Chabauty* *divisor*

$$D_{\mathrm{chab}} = \sum_{\widetilde{Q} \in \bar{X}(\mathbf{F}_{\mathfrak{p}})} n_{\widetilde{Q}}(\widetilde{Q}).$$

Then D_{chab} and $K_{\bar{X}} - D_{\mathrm{chab}}$ are both equivalent to effective divisors, so by Clifford's inequality (applied to the smooth proper curve $\bar{X}$) we have $r(D_{\mathrm{chab}}) := h^0(D_{\mathrm{chab}}) - 1 \leq \frac{1}{2}\deg(D_{\mathrm{chab}})$. On the other hand, by the semicontinuity of h^0 under specialization we have $h^0(D_{\mathrm{chab}}) \geq \dim V_{\mathrm{chab}} \geq g - r$. Combining these inequalities gives

$$\sum_{\widetilde{Q} \in \bar{X}(\mathbf{F}_{\mathfrak{p}})} n_{\widetilde{Q}} \leq 2r$$

which leads to Stoll's refinement of Coleman's bound.

Lorenzini and Tucker observed that one can generalize Coleman's bound to the case of bad reduction as follows. Let $\mathfrak{X}$ be a proper regular model for X over $\mathcal{O}_{\mathfrak{p}}$ and note that points of $X(K)$ specialize to $\bar{\mathfrak{X}}^{\mathrm{sm}}(\mathbf{F}_{\mathfrak{p}})$. One obtains by a similar argument the bound

$$\#X(K) \leq \sum_{\widetilde{Q} \in \bar{\mathfrak{X}}^{\mathrm{sm}}(\mathbf{F}_{\mathfrak{p}})} \left(1 + n_{\widetilde{Q}}\right), \tag{9}$$

where $n_{\widetilde{Q}}$ is the minimum over all nonzero ω in V_{chab} of $\mathrm{ord}_{\widetilde{Q}}\widetilde{\omega}$; here $\widetilde{\omega}$ denotes the reduction of ω to the unique irreducible component of the special fiber of $\mathfrak{X}$ containing $\widetilde{Q}$ and $\dim V_{\mathrm{chab}} \geq g - r > 0$ as before. Choosing a nonzero $\omega \in V_{\mathrm{chab}}$ as in Coleman's bound, the fact that the relative dualizing sheaf for $\mathfrak{X}$ has degree $2g-2$ gives the Lorenzini-Tucker bound.

In order to combine the bounds of Stoll and Lorenzini-Tucker, we see that it is natural to form the Chabauty divisor

$$D_{\mathrm{chab}} = \sum_{\widetilde{Q} \in \bar{\mathfrak{X}}^{\mathrm{sm}}(\mathbf{F}_{\mathfrak{p}})} n_{\widetilde{Q}}(\widetilde{Q})$$

and try to prove, using some version of semicontinuity of h^0 and Clifford's inequality, that its degree is at most $2r$. This is the main technical innovation of Katz and Zureick-Brown, so we state it as a theorem:

Theorem 6.15 (Katz–Zureick-Brown). *The degree of D_{chab} is at most $2r$.*

Combining Theorem 6.15 with (9) yields Theorem 6.14. As noted by Katz and Zureick-Brown, if one makes a base change from $K_{\mathfrak{p}}$ to an extension field K' over which there is a regular semistable model $\mathfrak{X}'$ for X dominating the base change of $\mathfrak{X}$, then the corresponding Chabauty divisors satisfy $D'_{\mathrm{chab}} \geq D_{\mathrm{chab}}$. (Here D'_{chab} is defined relative to the K'-vector space $V'_{\mathrm{chab}} = V_{\mathrm{chab}} \otimes_K K'$; one does not want to look at the Mordell-Weil group of J over extensions of K.) In order to prove Theorem 6.15, we may therefore assume that $\mathfrak{X}$ is a regular semistable model for X (and also that the residue field of K' is algebraically closed).

Let $d = \deg(D_{\mathrm{chab}})$. We now explain how to prove that $d \leq 2r$ when $\mathfrak{X}$ is a semistable regular model using limit linear series on metrized complexes of curves. Our proof is more direct than the one given in [KZB]; in particular, we do not need to introduce an "avoidance variant" of Clifford's inequality as in [KZB].

Proof of Theorem 6.15. Let $s = \dim_{K'} V'_{\mathrm{chab}} - 1 \geq g - r - 1 \geq 0$. We can identify V'_{chab} with an $(s+1)$-dimensional space W of rational functions on X in the usual way by identifying $H^0(X, \Omega^1_X)$ with $L(K_X) = \{f \ : \ \mathrm{div}(f) + K_X \geq 0\}$ for a canonical divisor K_X on X. The divisor D_{chab} on $\tilde{\mathfrak{X}}^{\mathrm{sm}}$ defines in a natural way a divisor $\mathcal{D} = \{D_\Gamma, \{D_v\}\}$ of degree d on the metrized complex $\mathfrak{C}\mathfrak{X}$ associated to $\mathfrak{X}$: the divisor D_v is the restriction of D_{chab} to C_v and $D_\Gamma = \sum_{v \in V(G)} \deg(D_v)(v)$. We promote the divisor $K_{\mathfrak{C}\mathfrak{X}} - \mathcal{D}$ to a limit linear series $\mathcal{L} = (K_{\mathfrak{C}\mathfrak{X}} - \mathcal{D}, \{H_v\})$ on $\mathfrak{C}\mathfrak{X}$ by defining H_v to be the reduction of W to C_v for each $v \in V(G)$. By the definition of D_{chab}, each element of H_v vanishes to order at least $n_{\widetilde{Q}}$ at each point $\widetilde{Q}$ in $\mathrm{supp}(D_v)$. By Theorems 1.8 and 6.9, $\mathcal{L}$ is a limit $\mathfrak{g}^s_{2g-2-d}$ on $\mathfrak{C}\mathfrak{X}$. In particular, we must have $r_{\mathfrak{C}\mathfrak{X}}(K_{\mathfrak{C}\mathfrak{X}} - \mathcal{D}) \geq g - r - 1$. On the other hand, Clifford's inequality for metrized complexes (Theorem 3.4) implies that $r_{\mathfrak{C}\mathfrak{X}}(K_{\mathfrak{C}\mathfrak{X}} - \mathcal{D}) \leq \frac{1}{2}(2g-2-d)$. Combining these inequalities gives $d \leq 2r$ as desired. □

7. Example: Generic chains of genus-one curves

In this section, we apply the theory developed in the previous section to the example of a generic chain of genus-one curves, and provide an analysis of the Brill-Noether theory of the corresponding metrized complexes. As we will see, the description of ranks of divisors on such metrized complexes turn out to be quite similar to the example considered in Cools et. al. [CDPR] of a generic chain of genus-one metric graphs. In particular, we provide an interpretation of the lingering paths of [CDPR] in terms of the η-function introduced in Section 5. In addition, we show that every divisor of degree d and rank r on such a metrized complex can be completed to a limit g^r_d, which does not hold for metrized complexes in general.

In addition to providing a concrete illustration of how to work with limit linear series on metrized complexes in practice, we hope that our analysis of this example provides some insight into the relationship between the proofs given in Eisenbud-Harris [EH] and [CDPR] of the non-existence part of the Brill-Noether theorem in algebraic geometry. The upshot of our analysis is that a generic chain of genus-one curves is Brill-Noether general not only with respect to limit linear series (which is known, see [O11, W]) but also with respect to the (unrestricted) rank function, while the tree of rational curves with genus-one tails used by Eisenbud and Harris is only Brill-Noether general with respect to limit linear series. In this sense, the Brill-Noether theory of a generic chain of genus-one curves admits a more 'purely combinatorial' analysis: one can prove that such curves are Brill-Noether general by a purely graph-theoretic argument, while for the reducible curves treated by Eisenbud and Harris one has to carefully study gluing conditions for certain spaces of rational functions.

7.1. **Rank of divisors on a generic chain of genus-one curves.** In Section 5.1.1 we provided a formula for the rank of divisors on a metric graph $\Gamma = \Gamma_1 \vee \Gamma_2$ obtained as a connected sum of two metric graphs Γ_1 and Γ_2. The discussion there can be generalized in a straightforward way to the context of metrized complexes as follows. Suppose the underlying graph of a metrized complex $\mathfrak{C}$ has a bridge edge e, and denote by $\mathfrak{C}_1$ and $\mathfrak{C}_2$ the two metrized complexes obtained by removing the edge e from $\mathfrak{C}$. Denote by x_1 and x_2 the points of $\mathfrak{C}_1$ and $\mathfrak{C}_2$, respectively, corresponding to the edge e.

Consider now two divisors $\mathcal{D}_1$ and $\mathcal{D}_2$ on $\mathfrak{C}_1$ and $\mathfrak{C}_2$, respectively, and denote by $\mathcal{D}$ the divisor on $\mathfrak{C}$ whose restriction to each $\mathfrak{C}_i$ is precisely $\mathcal{D}_i$. For a non-negative integer s, define

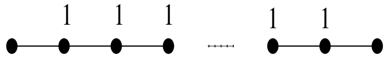

Figure 2. The dual graph of a chain of genus-one curves.

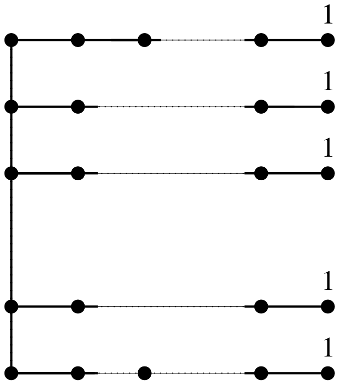

Figure 3. The dual graph of a tree of rational curves with genus-one tails.

$\eta(s)$ to be the smallest integer k such that $r_{\mathfrak{C}_1}(\mathcal{D}_1 + k(x_1)) = s$. Then

$$r_{\mathfrak{C}}(\mathcal{D}) = \min_{s\geq 0}\{s + r_{\mathfrak{C}_2}(\mathcal{D}_2 - \eta(s)(x_2))\}. \tag{10}$$

Formula (10) is very handy for dealing with metrized complexes $\mathfrak{C}$ whose underlying graph is a path, and essentially allows one to determine the rank of an arbitrary divisor on $\mathfrak{C}$ in a recursive way. Since our primary purpose here is to study chains of genus-one curves, we will restrict to this case, but it is not difficult to generalize the analysis that follows to chains of curves of arbitrary genus.

Consider a metrized complex $\mathfrak{C}$ whose dual graph G is the vertex-weighted graph with vertices $v_0, v_1, \dots, v_n, v_{n+1}$ of genus 0, 1, 1, …, 1, 0, respectively, depicted in Figure 2. The edge lengths are immaterial since G has genus zero; we may assume without loss of generality that every edge has length one. Denote by $C_0, C_1, \dots, C_n$, C_{n+1} the curves corresponding to the vertices $v_0, v_1, \dots, v_n, v_{n+1}$, where C_0 and C_{n+1} are rational curves and the other C_i are of genus one. At some point later we will suppose that for any genus-one curve C_i in $\mathfrak{C}$, $1 \leq i \leq n$, the marked points x_i, y_i on C_i corresponding to the two edges of G adjacent to v_i (namely $\{v_i, v_{i-1}\}$ and $\{v_i, v_{i+1}\}$, respectively) are *generic*. We will refer to such $\mathfrak{C}$ as a *generic chain of genus-one curves* (though it has two rational curves as tails; the rational tails make no difference from a purely mathematical point of view but they help to simplify the presentation of our arguments).

Consider now a divisor $\mathcal{D} = D_\Gamma \oplus \sum_{i=0}^{n+1} D_i$ of degree d on $\mathfrak{C}X$. We are interested in determining how large the rank of $\mathcal{D}$ can be. Since the answer depends only on the linear equivalence class of $\mathcal{D}$, we can and will assume that $\mathcal{D}$ is v_0-reduced. Since the underlying metric graph of $\mathfrak{C}X$ is a path and C_{n+1} has genus zero, it is easy to see that D_Γ is entirely supported on the vertices of Γ and $D_{n+1} = 0$.

Consider the edge $e_i = \{v_i, v_{i+1}\}$ in Γ, and denote by $\mathfrak{C}_{i,1}$ and $\mathfrak{C}_{i,2}$ the two metrized complexes obtained after removing e_i, which contain y_i and x_{i+1}, respectively. Denote by $\mathcal{D}_{i,1}$ and $\mathcal{D}_{i,2}$ the restriction of $\mathcal{D}$ to $\mathfrak{C}_{i,1}$ and $\mathfrak{C}_{i,2}$ respectively, and let $\eta_i : \mathbb{N} \cup \{0\} \to \mathbb{Z}$ be the function defined above for the bridge edge e_i in Γ. (Note that $\mathcal{D}_{i,1}$ and $\mathcal{D}_{i,2}$ are well-defined since by our assumption D_Γ is entirely supported on the vertices of Γ.) We have the following relation coming from (10):

$$r_{\mathfrak{C}X}(\mathcal{D}) = \min_{s\geq 0}\{s + r_{\mathfrak{C}_{i,2}}(\mathcal{D}_{i,2} - \eta_i(s)(x_{i+1}))\}.$$

It follows that the rank of $\mathcal{D}$ is determined as soon as the functions η_i are determined. In addition, the functions η_i satisfy similar recursive equations, as we will show in just a moment. We first note that η_0 has a particularly simple description. Denote by d_0 the coefficient of D_Γ at v_0. Since the curve corresponding to v_0 is rational, we have $\eta_0(s) = s - d_0$ for any $s \geq 0$.

Suppose now that we already know the function $\eta_i : \mathbb{N} \cup \{0\} \to \mathbb{Z}$. To determine η_{i+1}, we consider in the metrized complex $\mathfrak{C}_{i+1,1}$ the bridge edge e_{i+1} whose removal gives the metrized complex $\mathfrak{C}_{i,1}$ and the curve C_{i+1} seen as a metrized complex on a single vertex. By definition, $\eta_{i+1}(s)$ is the smallest integer satisfying $r_{\mathfrak{C}_{i+1,1}}(\mathcal{D}_{i+1,1} + \eta_{i+1}(s)(y_{i+1})) = s$. The recursive relation satisfied by the left-hand side of this equation gives

$$s = r_{\mathfrak{C}_{i+1,1}}(\mathcal{D}_{i+1,1}+\eta_{i+1}(s)(y_{i+1})) = \min_{t\geq 0}\big\{t+r_{C_{i+1}}(D_{i+1}+\eta_{i+1}(s)(y_{i+1})-\eta_i(t)(x_{i+1}))\big\}. \tag{11}$$

It follows that the function η_{i+1} can be calculated from the values of η_i and the rank function on C_{i+1}. As a consequence, once we know η_n we can determine the rank of $\mathcal{D}$.

7.1.1. *Brill-Noether theory on a generic chain of genus-one curves.* Consider now a divisor $\mathcal{D} = D_\Gamma \oplus \sum_{i=0}^{n} D_i$ on $\mathfrak{C}$ of degree d, that we suppose to be v_0-reduced. We may and do suppose in addition that D_i is effective for each $i \geq 1$. This easily implies:

- For $i \geq 1$, each D_i is of degree at most one.
- If $\deg(D_i) = 1$, then $D_i = (z_i)$ for a point z_i on $C_i \setminus \{x_i\}$.

We now present a criterion for the rank $r_{\mathfrak{C}}(\mathcal{D})$ to be at least r. Note that in taking the minimum in (11), the values of any function η_i over $t \geq r+1$ are automatically larger than r. In addition, the recursive equation relating η_{i+1} to η_i does not involve the value of η_i on larger integers. This means that if we are just interested in knowing whether or not $r_{\mathfrak{C}}(\mathcal{D}) \geq r$, we can restrict all functions η_i to the set $\{0, \dots, r\}$ and consider the values $\eta_i(0), \dots, \eta_i(r)$.

The following proposition summarizes the basic properties of η_i, and gives a necessary and sufficient condition for $r_{\mathfrak{C}}(\mathcal{D}) \geq r$ in terms of the values of $\eta_i(r)$.

Proposition 7.1. *Let r be a non-negative integer. Then:*

(i) *For $i = 0$, we have $\eta_0(s) = s - d_0$ for all s.*
(ii) *For each i, we have $\eta_i(0) < \eta_i(1) < \cdots < \eta_i(r)$.*
(iii) *$r_{\mathfrak{C}}(\mathcal{D}) \geq r$ if any only if for any $i = 0, \dots, n$ we have $\eta_i(r) \leq 0$.*

Proof. We have already proved (i). Statement (ii) is a direct consequence of the definitions of the rank function and η_i: we have

$$r_{\mathfrak{C}_{i,1}}(\mathcal{D}_{i,1} + (\eta_i(s+1)-1)(y_i)) \geq r_{\mathfrak{C}_{i,1}}(\mathcal{D}_{i,1} + \eta_i(s+1)(y_i)) - 1 = s$$

and therefore $\eta_i(s) \leq \eta_i(s+1) - 1$.

The "only if" direction in (iii) is a simple consequence of our assumption that $\mathcal{D}$ is v_0-reduced. Indeed, for each i, the divisor $-\eta_i(s)(x_{i+1}) + \mathcal{D}_{i,2}$ is v_{i+1}-reduced. Thus $r_{\mathfrak{C}_{i,2}}(-\eta_i(s)(x_{i+1}) + \mathcal{D}_{i,2}) \geq 0$ if and only if $-\eta_i(s)(x_{i+1}) + (z_{i+1})$ has non-negative rank on C_{i+1}. Since $x_{i+1} \neq z_{i+1}$, this is possible only if $\eta_i(s) \leq 0$.

To prove the other direction, we use the formula

$$r_{\mathfrak{C}X}(\mathcal{D}) = \min_{s\geq 0}\{s + r_{C_{n+1}}(D_{n+1} - \eta_n(s)(x_{n+1}))\}.$$

Since $C_{n+1} \simeq \mathbb{P}^1$ and

$$r_{\mathbb{P}^1}(D_{n+1} - \eta_n(s)(x_{n+1})) \geq -\eta_n(s) \geq -\eta_n(r) + r - s \geq r - s,$$

we obtain $r_{\mathfrak{C}X}(\mathcal{D}) \geq r$. □

We now use the recursive equation (11) to relate the values of η_{i+1} to the values of η_i. For fixed i, (11) tells us that

$$r_{C_{i+1}}(D_{i+1} + \eta_{i+1}(s)(y_{i+1}) - \eta_i(t)(x_{i+1})) \geq s - t$$

for every $t \leq s$, with equality for some value of t. Moreover, the inequalities for $t \leq s-2$ are implied by the inequality for $t = s-1$. Indeed, since $\eta_i(s-1) - \eta_i(t) \geq s-1-t$ and C_{i+1} is of genus one, the inequality for $s-1$ implies that

$$\begin{aligned} r_{C_{i+1}}(D_{i+1} + \eta_{i+1}(s)(y_{i+1}) - \eta_i(t)(x_{i+1})) &\geq r_{C_{i+1}}(D_{i+1} + \eta_{i+1}(s)(y_{i+1}) - \eta_i(s-1)(x_{i+1})) \\ &\qquad + s - t - 1 \\ &\geq s - t. \end{aligned}$$

Therefore the minimum in (11) is achieved for $t = s$ or $t = s-1$. For these two values of t we have

$$r_{C_{i+1}}(D_{i+1} + \eta_{i+1}(s)(y_{i+1}) - \eta_i(s)(x_{i+1})) \geq 0, \text{ and} \tag{12}$$

$$r_{C_{i+1}}(D_{i+1} + \eta_{i+1}(s)(y_{i+1}) - \eta_i(s-1)(x_{i+1})) \geq 1, \tag{13}$$

and $\eta_{i+1}(s)$ is defined in such a way that one of the two above inequalities is an equality.

The following cases can happen:

(a) $D_{i+1} = 0$.
(b) $D_{i+1} = (z_{i+1})$ for a point $z_{i+1} \in C_{i+1} \setminus \{x_{i+1}\}$.

In case (a), equations (12) and (13) tell us that $\eta_{i+1}(s) = \eta_i(s) + 1$ for any $0 \leq s \leq r$.

In case (b), since C_{i+1} is of genus one, we have $\eta_{i+1}(s) \in \{\eta_i(s), \eta_i(s) - 1\}$. In addition, $\eta_{i+1}(s) = \eta_i(s) - 1$ if and only if

(1) $(z_{i+1}) + (\eta_i(s) - 1)(y_{i+1}) - \eta_i(s)(x_{i+1}) \sim 0$ in C_{i+1} (by Equation (12)); and
(2) $\eta_i(s-1) \leq \eta_i(s) - 2$ (by Equation (13)).

Since x_{i+1} and y_{i+1} are generic points of C_{i+1}, relation (1) above can be satisfied by at most one value of $0 \leq s \leq r$. In other words, we have $\eta_{i+1}(t) = \eta_i(t)$ for all $0 \leq t \leq r$ except possibly for one value of s satisfying properties (1) and (2) above, for which we will have $\eta_{i+1}(s) = \eta_i(s) - 1$.

Consider now $\eta_0, \ldots, \eta_n$ as vectors of $\mathbb{Z}^{r+1}$ with basis $e_0, \ldots, e_r$ and define

$$\mathcal{A} := \{\alpha(0)e_0 + \cdots + \alpha(r)e_r \mid \alpha(0) < \cdots < \alpha(r) \leq 0\} \subset \mathbb{Z}^{r+1}.$$

Following [CDPR], a *lingering lattice path* in $\mathcal{A}$ of length n is a sequence of vectors $\alpha_0, \ldots, \alpha_n$ in $\mathcal{A}$ such that for each i exactly one of the following holds:

- $\alpha_{i+1} = \alpha_i + \sum_{s=0}^{r} e_s$
- $\alpha_{i+1} = \alpha_i - e_s$ for some $0 \leq s \leq r$
- $\alpha_{i+1} = \alpha_i$.

A lingering lattice path is called *of type* (d, d_0) if $\alpha_0 = \sum_s (s - d_0)e_s$ and the number of i such that $\alpha_{i+1} = \alpha_i - e_s$ is $d - d_0$. By the above discussion, each v_0-reduced divisor $\mathcal{D}$ of degree d and rank at least r defines a lingering path $\eta_0, \ldots, \eta_n$ in $\mathcal{A}$ of type (d, d_0) with $d = \deg(\mathcal{D})$ and $d_0 = \deg(D_0)$.

The following theorem, which is strongly reminiscent of the main results of [CDPR], summarizes the Brill-Noether theory of a generic chain $\mathfrak{C}$ of genus-one curves. Denote by $\mathcal{W}_d^r(\mathfrak{C})$ the locus of divisor classes of degree d and rank at least r in $\mathrm{Pic}_d(\mathfrak{C})$.

Theorem 7.2. *The map* $\mathcal{W}_d^r(\mathfrak{C}) \to \mathcal{A}$ *sending a divisor* $\mathcal{D}$ *to the lingering lattice path associated to the unique* v_0*-reduced divisor linearly equivalent to* $\mathcal{D}$ *is a surjection onto the union of the family of lingering paths in* $\mathcal{A}$ *of type* (d, d_0) *for some* $r \leq d_0 \leq d$*. In addition, the fiber over a lingering path* $\alpha_0, \ldots, \alpha_n$ *of type* (d, d_0) *is isomorphic to the product of the punctured curves* $C_i \setminus (\{x_i\} \cup B_i)$ *over all indices* i *such that* $\alpha_i = \alpha_{i-1}$*, where* B_i *is a finite set of points in* C_i *determined by* α_i*.*

Proof. The proof is almost implicit in the discussion preceding the theorem. Note that a divisor $\mathcal{D}$ corresponding to a lingering path of type (d, d_0) must have D_0 of degree d_0, and this determined D_0 up to linear equivalence since $C_0 \sim \mathbb{P}^1$. Moreover, the divisors D_{i+1} for $i \geq 0$ are completely determined as follows:

- If $\alpha_{i+1} = \alpha_i + \sum_s e_s$, then $D_{i+1} = 0$.
- If $\alpha_{i+1} = \alpha_i - e_s$ for some s, then $D_{i+1} = (z_{i+1})$ for the unique point z_{i+1} in $C_{i+1} \setminus \{x_{i+1}\}$ with the property that $(z_{i+1}) + (\eta_i(s) - 1)(y_{i+1}) - \eta_i(s)(x_{i+1}) \sim 0$ in C_{i+1}.
- If $\alpha_{i+1} = \alpha_i$, then D_{i+1} can be (z_{i+1}) for any point z_{i+1} of $C_{i+1} \setminus \{x_{i+1}\}$ which does not belong to the finite subset $B_{i+1} \subset C_{i+1}$ consisting of those points z with $(z) + (\alpha_i(s) - 1)(y_{i+1}) - \alpha_i(s)(x_{i+1}) \sim 0$ in C_{i+1} for some integer $0 \leq s \leq r$ with $\alpha(s-1) \leq \alpha(s) - 2$.
- $D_{n+1} = 0$.

Since the lingering lattice path we started with is of type (d, d_0), the divisor $\mathcal{D}$ constructed in this way will have degree d. It is easy to see that $\mathcal{D}$ is v_0-reduced. Moreover, by the definition of η_i we have $\eta_i = \alpha_i$ for all $i \leq n$, which shows that $\mathcal{D} \in \mathcal{W}_d^r(\mathfrak{C})$, proving surjectivity. The explicit description of the fibers of the map $\mathcal{W}_d^r \to \mathcal{A}$ follows from our discussion of the possible choices for the points z_i in the case where $\alpha_i = \alpha_{i-1}$. □

Corollary 7.3. *Let* $\rho = n - (r+1)(n - d + r)$*. Then for any lingering path* $\alpha = \alpha_0, \ldots, \alpha_n$ *of type* (d, d_0)*, the dimension of the fiber over* α *in the map* $\mathcal{W}_d^r \to \mathcal{A}$ *is at most* $\rho + r - d_0$*.*

Proof. First note that by Theorem 7.2, the dimension of the fiber in question is the number of indices i with $\alpha_i = \alpha_{i-1}$, so we just need to bound the number of such indices. The argument is similar to the reasoning in [CDPR]. Let $\mathcal{D}$ be a v_0-reduced divisor whose image is the lingering lattice path α. Since the degree of $\mathcal{D}$ is d, there are exactly $n - d + d_0$ indices i with $\alpha(i+1) = \alpha(i) + \sum_s e_s$. The coordinate $\alpha_n(r)$ is equal to $r - d_0 + n - d + d_0 - a_r$, where a_r is the number of indices i such that $\alpha_{i+1} = \alpha_i - e_r$. This shows that $a_r \geq r + n - d$. The number a_s of indices i such that $\alpha_{i+1} = \alpha_i - e_s$ is at least a_r, so the number of indices i with $\alpha_{i+1} = \alpha_i$ is at most $n - (r+1)a_r - n + d - d_0$. A simple calculation shows that this is at most $\rho + r - d_0$. □

Corollary 7.4. *If* $\rho = 0$*, then the map* $\mathcal{W}_d^r(\mathfrak{C}) \to \mathcal{A}$ *is a bijection from* $\mathcal{W}_d^r$ *to the set of lingering paths of type* (d, r)*. In addition,* $\eta_n(s) = s - r$ *for any* $s = 0, \ldots, r$*.*

Proof. Let $\mathcal{D}$ be a v_0-reduced divisor of degree d and rank r. Obviously $d_0 \geq r$. Since $\rho = 0$, the dimension of each fiber is at most $r - d_0$. Thus $r = d_0$, and by Theorem 7.2 we get a bijection between $\mathcal{W}_d^r$ and the set of lingering lattice paths of type (d, r) in $\mathcal{A}$.

For the second statement, note that by the proof of Corollary 7.3, when $\rho = 0$ the number a_s of indices i with $\alpha_{i+1} = \alpha_i - e_s$ is equal to $r + n - d$ for $s = 0, \ldots, r$. In addition, $d_0 = r$ implies that $\eta_0(s) = s - r$ for all s, and thus the number of indices i with $e_{i+1} = e_i + \sum_s \alpha_s$ is equal to $n - d + r$. A simple calculation now shows that $\eta_n(s) = \eta_0(s) = s - r$ for $s = 0, \ldots, r$. □

Finally, we show that every divisor class of degree d and rank r on $\mathfrak{C}$ can be completed to a limit $\mathfrak{g}_d^r$. The fact that there is no obstruction to doing this is a special property of the particular metrized complexes which we are considering in this section.

Proposition 7.5. *Let $\mathcal{D}$ be a divisor of degree d and rank r on $\mathfrak{C}$, which we may assume without loss of generality to be v_0-reduced. Then there exists a collection of κ-vector spaces $\mathcal{H} = \{H_0, \ldots, H_{n+1}\}$, with H_i an $(r+1)$-dimensional subspace of $\kappa(C_i)$ for all i, such that the pair $(\mathcal{D}, \mathcal{H})$ represents a limit linear $\mathfrak{g}_d^r$ on $\mathfrak{C}$.*

Proof. We begin by defining the collection $\mathcal{H}$. Since C_0 is rational, we may assume by Corollary 7.4 that $D_0 = d_0(y_0) = r(y_0)$. For $i = 0$, define H_0 to be the $(r+1)$-dimensional vector space of all rational functions f on C_0 with $r(y_0) + \mathrm{div}(f) \geq 0$. For any $1 \leq i \leq n$, Equation (12) implies that for each s there is a rational function $f_{i,s}$ on C_i such that

$$D_i + \eta_i(s)(y_i) - \eta_{i-1}(s)(x_i) + \mathrm{div}(f_{i,s}) \geq 0.$$

Define $H_i = \{f_{i,0}, \ldots, f_{i,r}\}$. Similarly, define H_{n+1} to be the $(r+1)$-dimensional space of all rational functions f on $C_{n+1} \cong \mathbb{P}^1$ with $r(x_{n+1}) + \mathrm{div}(f) \geq 0$.

We need to prove that $r_{\mathfrak{C},\mathcal{H}}(\mathcal{D}) = r$. For this, it suffices to show that for each curve C_i, and for any $s_1 \leq s_2$, the intersection of H_i with $H^0\Big(C_i, \mathcal{O}\big(D_i - \eta_{i-1}(s_1)(x_i) + \eta_i(s_2)(y_i)\big)\Big)$ has dimension at least $s_2 - s_1 + 1$. (Here $\eta_{-1}(.) = \eta_{n+1}(.) = 0$.) This follows from the fact that all of the functions $f_{i,s}$ for $s = s_1, \ldots, s_2$ belong to this intersection; indeed, for such s we have

$$-\eta_{i-1}(s)(x_i) + \eta_i(s)(y_i) \leq -\eta_{i-1}(s_1)(x_i) + \eta_i(s_2)(y_i).$$

□

Appendix A. Rank-determining sets for metrized complexes

We retain the terminology from Section 2. Let $\mathfrak{C}$ be a metrized complex of algebraic curves, Γ the underlying metric graph, $G = (V, E)$ the model of Γ and $\{C_v\}_{v \in V}$ the collection of smooth projective curves over κ corresponding to $\mathfrak{C}$. In this section, we generalize some basic results concerning rank-determining sets [L, HKN] from metric graphs to metrized complexes.

Let R be a subset of the union of $C_v(\kappa)$ for $v \in V$. An effective divisor $\mathcal{E} = E_\Gamma \oplus \sum_v E_v$ on $\mathfrak{C}$ is said to have *support in* R if $\mathrm{supp}(E_\Gamma) \subseteq V$ and $\mathrm{supp}(E_v) \subseteq R$.

A subset $R \subset \bigcup_{v \in V} C_v(\kappa)$ is called *rank-determining* if for any divisor $\mathcal{D}$ on $\mathfrak{C}$, $r_{\mathfrak{C}}(\mathcal{D})$ coincides with the restricted rank $r_{\mathfrak{C}}^R(\mathcal{D})$, defined as the largest integer k such that $\mathcal{D} - \mathcal{E}$ is linearly equivalent to an effective divisor for all degree k effective divisors $\mathcal{E}$ on $\mathfrak{C}$ supported in R. In other words, R is rank-determining if in the definition of rank given in Section 2, one can restrict to effective divisors $\mathcal{E}$ with support in R.

The following theorem is a common generalization of (a) Luo's theorem [L] (see also [HKN]) that V is a rank-determining set for any loopless model $G = (V, E)$ of a metric graph Γ and (b) the classical fact (also reproved in [L]) that for any smooth projective curve C of genus g over κ, every subset of $C(\kappa)$ of size $g+1$ is rank-determining.

Theorem A.1. *Let $\mathfrak{C}$ be a metrized complex of algebraic curves, and suppose that the given model G of Γ is loopless. Let $R_v \subset C_v(\kappa)$ be a subset of size $g_v + 1$ and $R = \cup_{v\in V} R_v$. Then R is a rank-determining subset of $\mathfrak{C}$.*

Let $\mathcal{D}$ be a divisor on $\mathfrak{C}$. For any point $P \in \Gamma$, denote by $\mathcal{D}^P = D^P_\Gamma \oplus \sum_{v\in V} D^P_v$ the quasi-unique P-reduced divisor on $\mathfrak{C}$ linearly equivalent to $\mathcal{D}$.

Lemma A.2. *A divisor $\mathcal{D}$ on $\mathfrak{C}$ has rank at least one if and only if*

1. *For any point P of Γ, $D^P_\Gamma(P) \geq 1$, and*
2. *For any vertex $v \in V(G)$, the divisor D^v_v has rank at least one on C_v.*

Proof. The condition $r_{\mathfrak{C}}(\mathcal{D}) \geq 1$ is equivalent to $r_{\mathfrak{C}}(\mathcal{D} - \mathcal{E}) \geq 0$ for any effective divisor $\mathcal{E}$ of degree one on $\mathfrak{C}$. For any point P in $\Gamma \setminus V$, $\mathcal{D} - (P)$ has non-negative rank if and only if $D^P_\Gamma(P) - 1 \geq 0$ (by Lemma 3.10.) Similarly, for any $v \in V$ and $x \in C_v(\kappa)$, the divisor $\mathcal{D} - (v) \oplus (x)$ has non-negative rank in $\mathfrak{C}$ if and only if $D^v_\Gamma(v) \geq 1$ and $D^v_v - (x)$ has non-negative rank on C_v (by Lemma 3.10). These are clearly equivalent to (1) and (2), and the lemma follows. □

Lemma A.3. *A subset $R \subseteq \bigcup_{v\in V} C_v(\kappa)$ which has non-empty intersection with each $C_v(\kappa)$ is rank-determining if and only if for every divisor $\mathcal{D}$ of non-negative rank on $\mathfrak{C}$, the following two assertions are equivalent:*

(i) $r_{\mathfrak{C}}(\mathcal{D}) \geq 1$.
(ii) *For any vertex $u \in V$, and for any point $z \in R \cap C_u(\kappa)$, $D^u_u - (z)$ has non-negative rank on C_u*

Proof. In view of Lemma A.2, for a rank-determining set the two conditions (i) and (ii) are equivalent. Suppose now that (i) and (ii) are equivalent for any divisor $\mathcal{D}$ on $\mathfrak{C}$. By induction on r, we prove that $r_{\mathfrak{C}}(\mathcal{D}) \geq r$ if and only if for every effective divisor $\mathcal{E}$ of degree r with support in R, $r_{\mathfrak{C}}(\mathcal{D} - \mathcal{E}) \geq 0$. This will prove that R is rank-determining.

The case $r = 1$ follows by the hypothesis and Lemma A.2. Supposing now that the statement holds for some integer $r \geq 1$, we prove that it also holds for $r + 1$. Let $\mathcal{D}$ be a divisor with the property that for any effective $\mathcal{E}$ of degree $r+1$ with support in R, $r_{\mathfrak{C}}(\mathcal{D} - \mathcal{E}) \geq 0$. Fix an effective divisor $\mathcal{E}$ of degree r with support in R. By the base of our induction, the divisor $\mathcal{D} - \mathcal{E}$ has rank at least one on $\mathfrak{C}$ because for any effective divisor $\mathcal{E}'$ of degree one with support in R, $r_{\mathfrak{C}}(\mathcal{D} - \mathcal{E} - \mathcal{E}') \geq 0$. Thus, for any divisor $\mathcal{E}'$ of degree one on $\mathfrak{C}$, $r_{\mathfrak{C}}(\mathcal{D} - \mathcal{E}' - \mathcal{E}) \geq 0$. This holds for any effective divisor $\mathcal{E}$ of degree r with support in R. From the inductive hypothesis, we infer that $\mathcal{D} - \mathcal{E}'$ has rank at least r on $\mathfrak{C}$. Since this holds for any effective divisor $\mathcal{E}'$ of degree one on $\mathfrak{C}$, we conclude that $\mathcal{D}$ has rank at least $r + 1$. □

Let $\mathcal{D} = D_\Gamma \oplus \sum_v D_v$ be a divisor of degree d and non-negative rank on $\mathfrak{C}$. Define

$$|D_\Gamma| := \{E \geq 0 \,|\, E \in \mathrm{Div}(\Gamma) \text{ and } D_\Gamma \sim E\}.$$

Note that $|D_\Gamma|$ is a non-empty subset of the symmetric product $\Gamma^{(d)}$ of d copies of Γ.

Consider the *reduced divisor* map $\mathrm{Red}^{\mathfrak{C}}_{\mathcal{D}} : \Gamma \to \Gamma^{(d)}$ which sends a point $P \in \Gamma$ to D^P_Γ. (Recall that $\mathcal{D}^P = D^P_\Gamma \oplus \sum_v D^P_v$ is the quasi-unique P-reduced divisor associated to $\mathcal{D}$ in $\mathfrak{C}$.) The following theorem extends [A, Theorem 3] to divisors on metrized complexes.

Lemma A.4. *For any divisor $\mathcal{D} = D_\Gamma \oplus \Sigma_v D_v$ of degree d and non-negative rank on $\mathfrak{C}$, the reduced divisor map $\mathrm{Red}^{\mathfrak{C}}_{\mathcal{D}} : \Gamma \to \Gamma^{(d)}$ is continuous.*

Proof. As with the corresponding result in [A], the proof is based on an explicit description of the infinitesimal behavior of the reduced divisor map $\mathrm{Red}^{\mathfrak{C}}_{\mathcal{D}}$ around any point of Γ.

Let P be a point of Γ and let $\vec{\mu}$ be a tangent direction in Γ emanating from P. For $\epsilon > 0$ sufficiently small, we denote by $P + \epsilon\vec{\mu}$ the point of Γ at distance ϵ from P in the direction of $\vec{\mu}$. Let $\mathcal{D}^P = D^P_\Gamma \oplus \sum_{v\in V} D^P_v$ be the P-reduced divisor linearly equivalent to $\mathcal{D}$ in $\mathfrak{C}$. Then either

1. $\mathcal{D}^P$ is $(P + \epsilon\vec{\mu})$-reduced for all sufficiently small $\epsilon > 0$; or
2. There is a (maximal) saturated cut S in Γ with respect to $\mathcal{D}^P$ such that $P \in \partial S$ and $P + \epsilon\vec{\mu} \notin S$ for all sufficiently small $\epsilon > 0$ (see Section 3.2 for the definition of 'saturated').

In the first case, $\mathrm{Red}^{\mathfrak{C}}_{\mathcal{D}}$ is constant on the segment $[P, P + \epsilon\vec{\mu}]$ for all sufficiently small $\epsilon > 0$. In the second case, there exists an $\epsilon_0 > 0$ such that for any $0 < \epsilon < \epsilon_0$, the reduced divisor $\mathcal{D}^{P+\epsilon\vec{\mu}} = D^{P+\epsilon\vec{\mu}}_\Gamma \oplus \sum_{v\in V} D^{P+\epsilon\vec{\mu}}_v$ has the following description.

For any sufficiently small $\gamma > 0$ and any positive integer $\alpha > 0$, we will define below a rational function $f^{(\gamma,\alpha)}_\Gamma$ on Γ. For an appropriate choice of $\gamma = \gamma(\epsilon)$ and $\alpha = \alpha(\epsilon)$, we will have $\mathcal{D}^{P+\epsilon\vec{\mu}} = \mathcal{D}^P + \mathrm{div}(\mathfrak{f}^{\gamma,\alpha})$, where $\mathfrak{f}^{\gamma,\alpha}$ is the rational function on $\mathfrak{C}$ given by $f^{\gamma,\alpha}_\Gamma$ on Γ and $f_v = 1$ on each C_v.

Let $\vec{\mu}_1, \dots, \vec{\mu}_s$ different from $\vec{\nu}$ be all the different out-going tangent vectors to Γ at the boundary points $x_1, \dots, x_s \in \partial S$, respectively. (In particular, it might be the case that $x_i = x_j$ for two different indices i and j). Let $\gamma > 0$ be small enough so that for any point $x \in \partial S$ and any out-going unit vector $\vec{\mu}^*$ tangent to Γ at x, the whole segment $(x, x + \gamma\vec{\mu}^*\,]$ lies outside S and does not contain any point of the support of D_Γ. The intervals $(x_i, x_i + \gamma\vec{\mu}_i\,]$ together with $[P, P + \frac{\gamma}{\alpha}\vec{\mu}\,]$ form (a segment on) all the out-going edges from S.

The rational function $f^{(\gamma,\alpha)}_\Gamma$ on Γ is defined as follows.

- The restriction of $f^{(\gamma,\alpha)}_\Gamma$ to S is zero;
- The restriction of $f^{(\gamma,\alpha)}_\Gamma$ to any interval $[x_i, x_i + \gamma\vec{\mu}_i]$ is linear of slope -1 for any $1 \le i \le s$;
- The restriction of $f^{(\gamma,\alpha)}_\Gamma$ to the interval $[P, P + (\frac{\gamma}{\alpha})\vec{\mu}\,]$ is linear of slope $-\alpha$; and
- $f^{(\gamma,\alpha)}_\Gamma$ takes value $-\gamma$ at any other point of Γ.

It remains to determine the values of γ and α as functions of ϵ. We will explain in a moment how to determine the value of α. Once this has been done, the value of γ will be equal to $\alpha\epsilon$, so that the point $P + (\frac{\gamma}{\alpha})\vec{\mu}$ coincides with the point $P + \epsilon\vec{\mu}$. It remains to determine α. Consider the following two cases depending on whether or not P is a vertex of G:

- If $P \in \Gamma \setminus V$, then $\alpha = D^P_\Gamma(P) - \mathrm{outdeg}_S(P) + 1$. Since S is a saturated cut in Γ for D^P_Γ, we have $D^P_\Gamma(P) \ge \mathrm{outdeg}_S(P)$, and thus $\alpha \ge 1$.

- If $P = v$ for a vertex $v \in V(G)$, let $e_0, e_1, \dots, e_l$ be the out-going edges at v with respect to S, and consider the points $x_v^{e_0}, x_v^{e_1}, \dots, x_v^{e_l}$ in $C_v(\kappa)$ indexed by these edges. Suppose in addition that e_0 is the edge which corresponds to the tangent direction μ. Since S is a saturated cut with respect to D_Γ^v, the divisor $D_v - \mathrm{div}_v(\partial S) = D_v - \sum_{i=0}^{l}(x_v^{e_i})$ has non-negative rank in C_v. Define α to be the largest integer n such that $D_v - n(x_v^{e_0}) - \sum_{i=1}^{l}(x_v^{e_i})$ has non-negative rank. Obviously, $\alpha \geq 1$.

□

We are now ready to give the proof of Theorem A.1.

Proof of Theorem A.1. By Lemma A.3, it is enough to check the equivalence of the following two properties for any divisor $\mathcal{D}$ on $\mathfrak{C}$:

(i) $r_{\mathfrak{C}}(\mathcal{D}) \geq 1$.
(ii) For any $u \in V$ and any point $z \in R_u = R \cap C_u(\kappa)$, the divisor $D_u^u - (z)$ has non-negative rank on C_u.

It is clear that (i) implies (ii). So we only need to prove that (ii) implies (i). In addition, by Lemma A.2, Property (i) is equivalent to

(1) for any point P of Γ, $D_\Gamma^P(P) \geq 1$; and
(2) for any vertex $v \in V(G)$, the divisor D_v^v has rank at least one on C_v.

So it suffices to prove (ii) $\Rightarrow$ (1) $\wedge$ (2). Since cardinality of R_v is $g_v + 1$, R_v is rank-determining in C_v. Therefore, (ii) implies (2). We now show that (2) implies (1). Let Γ_0 be the set of all $P \in \Gamma$ such that $D_\Gamma^P(P) \geq 1$. By the continuity of the map $\mathrm{Red}_{\mathcal{D}}^{\mathfrak{C}}$, Γ_0 is a closed subset of Γ. In addition, since D_v^v has rank at least one on C_v and $D_\Gamma^v(v) = \deg(D_v^v)$ for every vertex $v \in C$, we have $V \subset \Gamma_0$. This shows that $\Gamma \setminus \Gamma_0$ is a disjoint union of open segments contained in edges of G. Suppose for the sake of contradiction that $\Gamma_0 \subsetneq \Gamma$, and let $I = (P, Q)$ be a non-empty segment contained in the edge $\{u, v\}$ of G such that $I \cap \Gamma_0 = \emptyset$. We first observe that by the above description of the reduced divisor map, $\mathrm{Red}_{\mathcal{D}}^{\mathfrak{C}}$ must be constant on the closed interval $[P, Q]$. To see this, note that for any point $Z \in [P, Q]$ and any tangent direction $\vec{\mu}$ for which $Z + \epsilon\vec{\nu} \in [P, Q]$ whenever $\epsilon > 0$ is sufficiently small, we are always in the case (1) in the description of $\mathrm{Red}_{\mathcal{D}}^{\mathfrak{C}}$. Otherwise, there would be an integer $\alpha > 0$ such that $\mathcal{D}^{Z+\epsilon\vec{\mu}} = \mathcal{D}^Z + \mathrm{div}(\mathfrak{f}^{(\epsilon\alpha,\alpha)})$ for all sufficiently small $\epsilon > 0$. In particular, this would imply that $D_\Gamma^{Z+\epsilon\vec{\mu}}(Z + \epsilon\vec{\mu}) = \alpha \geq 1$, by the definition of $f^{(\eta,a)}$, which implies that $Z + \epsilon\vec{\nu} \in \Gamma_0$, a contradiction. This shows that $\mathrm{Red}_{\mathcal{D}}^{\mathfrak{C}}$ is constant on $[P, Q]$.

A case analysis (depending on whether P or Q are vertices of the graph or not) shows that for a point $Z \in (P, Q)$, the cut $S = \Gamma \setminus (P, Q)$ is saturated for $\mathcal{D}^P = \mathcal{D}^Q$. Since $\mathcal{D}^Z = \mathcal{D}^P = \mathcal{D}^Q$, and S does not contain Z, this contradicts the assumption that $\mathcal{D}^Z$ is Z-reduced. □

Theorem A.1 has the following direct corollaries.

Corollary A.5. *Let $\mathfrak{C}$ be a metrized complex with underlying metric graph Γ and with model $G = (V, E)$. Let $\mathcal{G}$ be a subgroup of $\mathbb{R}$ which contains all the edge lengths in G. For any divisor $\mathcal{D} \in \mathrm{Div}(\mathfrak{C})_{\mathcal{G}}$, we have*

$$r_{\mathfrak{C},\mathcal{G}}(\mathcal{D}) = r_{\mathfrak{C}}(\mathcal{D}).$$

Proof. Fix a rank-determining set $R \subset \cup_{v\in V} C_v(\kappa)$ as in Theorem A.1. Since R is rank-determining and any effective divisor $\mathcal{E}$ with support in R obviously belongs to $\mathrm{Div}(\mathfrak{C})_{\mathcal{G}}$, to prove the equality of $r_{\mathfrak{C},\mathcal{G}}(\mathcal{D})$ and $r_{\mathfrak{C}}(\mathcal{D})$ it will be enough to show that the two statements $r_{\mathfrak{C},\mathcal{G}}(\mathcal{D}) \geq 0$ and $r_{\mathfrak{C}}(\mathcal{D}) \geq 0$ are equivalent. Obviously, the former implies the latter, so we only need to show that if $r_{\mathfrak{C}}(\mathcal{D}) \geq 0$ then $r_{\mathfrak{C},\mathcal{G}}(\mathcal{D}) \geq 0$. Let v be a vertex of G and $\mathcal{D}^v$ the v-reduced divisor linearly equivalent to $\mathcal{D}$. By Lemma 3.10, $r_{\mathfrak{C}}(\mathcal{D}) \geq 0$ is equivalent to $r_{C_v}(D^v_v) \geq 0$. Now let $\mathcal{D}$ be an element of $\mathrm{Div}(\mathfrak{C})_{\mathcal{G}}$ with $r_{C_v}(D^v_v) \geq 0$. Since $v \in V$ and $\mathcal{G}$ contains all the edge-lengths in G, it is easy to see that $\mathcal{D}$ and $\mathcal{D}^v$ differ by the divisor of a rational function $\mathfrak{f}$ with support in $\mathrm{Div}(\mathfrak{C})_{\mathcal{G}}$. In other word, $\mathcal{D} \sim \mathcal{D}^v$ in $\mathrm{Div}(\mathfrak{C})_{\mathcal{G}}$. Since $\mathcal{D}^v$ is linearly equivalent to an effective divisor in $\mathrm{Div}(\mathfrak{C})_{\mathcal{G}}$ (with constant rational function on Γ), we conclude that $r_{\mathfrak{C},\mathcal{G}}(D) \geq 0$. ☐

Corollary A.6. *Let $\mathfrak{C}X_0$ be the regularization of a strongly semistable curve X_0 over κ. Let $\mathcal{L}$ be a line bundle on X_0 corresponding to a divisor $\mathcal{D} \in \mathrm{Div}(\mathfrak{C})$. Then $r_c(\mathcal{L}) = r_{\mathfrak{C}X_0}(\mathcal{D})$.*

Proof. This follows from the previous corollary with $\mathcal{G} = \mathbb{Z}$. ☐

References

[A] O. Amini, *Reduced divisors and embeddings of tropical curves.* Trans. Amer. Math. Soc., to appear.

[ABBR] O. Amini, M. Baker, E. Brugallé, and J. Rabinoff, *Lifting harmonic morphisms of tropical curves, metrized complexes, and Berkovich skeleta.* In preparation.

[AC] O. Amini and L. Caporaso, *Riemann-Roch theory for weighted graphs and tropical curves.* Preprint available at `arXiv:1112.5134`.

[B] M. Baker, *Specialization of linear systems from curves to graphs.* Algebra & Number Theory 2 (2008), no. 6, 613–653.

[BN] M. Baker and S. Norine, *Riemann-Roch and Abel-Jacobi theory on a finite graph.* Adv. Math. 215 (2007), no. 2, 766–788.

[BPR] M. Baker, S. Payne, and J. Rabinoff, *Non-Archimedean geometry, tropicalization, and metrics on curves.* Preprint available at `arXiv:1104.0320v1`.

[Ber] V. G. Berkovich, *Spectral theory and analytic geometry over non-Archimedean fields.* Mathematical Surveys and Monographs, vol. 33, American Mathematical Society, Providence, RI, 1990.

[Ber99] V. G. Berkovich, *Smooth p-adic analytic spaces are locally contractible.* Invent. Math. 137 (1999), no. 1, 1–84.

[Ber04] V. G. Berkovich, *Smooth p-adic analytic spaces are locally contractible II.* Geometric aspects of Dwork theory, Walter de Gruyter and Co. KG, Berlin, 2004, 293–370.

[Bi] M. T. I. Bigas, *Brill-Noether theory for stable vector bundles.* Duke Math. J. 62 (1991), no. 2 385–400.

[C1] L. Caporaso, *Linear series on semistable curves.* Int. Math. Res. Not. 13 (2011), 2921–2969.

[CDPR] F. Cools, J. Draisma, S. Payne, and E. Robeva, *A tropical proof of the Brill-Noether Theorem.* Adv. Math. 230 (2012), no. 2, 759–776.

[Co] R. F. Coleman, *Effective Chabauty,* Duke Math. J. 52 (1985), no. 3, 765–770.

[DM] P. Deligne and D. Mumford, *The irreducibility of the space of curves of given genus.* Publications Mathématiques de l'IHES 36 (1969), no. 1, 75–109.

[EH] D. Eisenbud and J. Harris, *Limit linear series: Basic theory.* Invent. Math. 85 (1986), 337-371.

[EH2] D. Eisenbud and J. Harris, *The Kodaira dimension of the moduli space of curves of genus* ≥ 23. Invent. Math. 90 (1987), no. 2, 359–387.

[EH3] D. Eisenbud and J. Harris, *Existence, decomposition, and limits of certain Weierstrass points.* Invent. Math. 87 (1987), no. 3, 495–515.

[EH4] D. Eisenbud and J. Harris, *The monodromy of Weierstrass points.* Invent. Math. 90 (1987), no.2, 333–341.

[E] E. Esteves, *Linear systems and ramification points on reducible nodal curves.* Mathematica Contemporanea 14 (1998), 21–35.

[EM] E. Esteves and N. Medeiros, *Limit canonical systems on curves with two components.* Invent. Math. 149 (2002), no. 2, 267–338.

[GK] A. Gathmann and M. Kerber, *A Riemann-Roch theorem in tropical geometry.* Math. Z. 259 (2008), no. 1, 217–230.

[HM98] J. Harris and I. Morrison, *Moduli of Curves.* Graduate Texts in Mathematics 187, Springer Verlag, 1998.

[HM82] J. Harris and D. Mumford, *On the Kodaira dimension of the moduli space of curves.* Invent. Math. 67 (1982), 23–88.

[H] R. Hartshorne, *Residues and Duality.* Springer Lecture Notes **20**, 1966.

[HMY] C. Hasse, G. Musiker, and J. Yu, *Linear systems on tropical curves*, Math. Z. 270 (2012), no. 3-4, 1111–1140.

[HKN] J. Hladký, D. Kràl', and S. Norine, *Rank of divisors on tropical curves.* Preprint available at `arXiv:0709.4485`.

[KZB] E. Katz and D. Zureick-Brown, *The Chabauty-Coleman bound at a prime of bad reduction and Clifford bounds for geometric rank functions.* Preprint available at `arxiv.org/1204.3335`.

[LT] D. J. Lorenzini and T. J. Tucker, *Thue equations and the method of Chabauty-Coleman*, Invent. Math. 148 (2002), no. 1, 47–77.

[LPP] C. M. Lim, S. Payne, and N. Potashnik, *A note on Brill-Noether thoery and rank determining sets for metric graphs.* Int. Math. Res. Not., to appear.

[L] Y. Luo, *Rank-determining sets of metric graphs.* J. Comb. Theory. Series A. 118 (2011), no. 6, 1775–1793.

[MP] W. McCallum and B. Poonen, *The method of Chabauty and Coleman*, June 14, 2010. Preprint, `http://www-math.mit.edu/~poonen/papers/chabauty.pdf`, to appear in Panoramas et Synthèses, Société Math. de France.

[MZ] G. Mikhalkin and I. Zharkov, *Tropical curves, their Jacobians and Theta functions.* Contemporary Mathematics 465: Proceedings of the International Conference on Curves and Abelian Varieties in honor of Roy Smith's 65th birthday. 203–231.

[O1] B. Osserman, *A limit linear series moduli scheme (Un schéma de modules de séries linéaires limites).* Ann. Inst. Fourier 56 (2006), no. 4, 1165–1205.

[O2] B. Osserman, *Linked Grassmannians and crude limit linear series.* Int. Math. Res. Not. 25 (2006), 1–27.

[O11] B. Osserman, *A simple characteristic-free proof of the Brill-Noether theorem.* http://arxiv.org/pdf/1108.4967v1.

[Par] B. Parker, *Exploded manifolds.* Adv. Math. 229 (2012), no. 6, 3256–3319.

[Pay] S. Payne, *Fibers of tropicalization.* Math. Zeit. 262 (2009), 301–311.

[R] Z. Ran, *Boundary modifications of Hodge bundles and enumerative geometry.* Preprint available at `arxiv.org/1011.0406`.

[St] M. Stoll, *Independence of rational points on twists of a given curve*, Compos. Math. 142 (2006), no. 5, 1201–1214.

[W] G. Welters, *A theorem of Gieseker-Petri type for Prym varieties.* Ann. Sci. Éc. Norm . Supér. 18 (1985), no. 4, 671–683.

CNRS - DMA, École Normale Supérieure, Paris
E-mail address: `oamini@math.ens.fr`

Department of Mathematics, University of California, Berkeley and School of Mathematics, Georgia Institute of Technology
E-mail address: `mbaker@math.gatech.edu`